



UNIVERSITY OF TRENTO

DEPARTMENT OF PHYSICS

DISSERTATION

~ · ~

ACADEMIC YEAR 2025/2026

Optimizing Quantum Annealing for Trapped Ions: Performance, Resources, and Noise Mitigation

Supervisors

Prof. Philipp HAUKE
Prof. Enrico BLANZIERI

PhD Candidate

Sebastian NAGIES

FINAL EXAMINATION DATE: June 25, 2026

Abstract

Quantum annealing is a promising method for solving combinatorial optimization problems with quantum computers, but its practical performance on near-term devices is constrained by the architecture of the underlying platform, the presence of noise, and the design of the annealing protocol itself. This thesis addresses these challenges along two complementary directions, using trapped-ion hardware based on the Magnetic Gradient Induced Coupling (MAGIC) scheme as a reference architecture.

On the hardware side, we analyze the MAGIC setup beyond its standard leading-order description, computing higher-order corrections arising from anharmonicities of the Coulomb repulsion and the trapping potential, as well as from magnetic-field curvature. Most resulting terms are negligible in realistic situations, with the notable exceptions of a phonon-occupation-dependent longitudinal field that grows with chain length and a two-to-one phonon conversion process. To address fluctuations of external magnetic fields, the dominant error channel for magnetically sensitive qubits, we discuss an active noise mitigation protocol for quantum annealing in which the optimization problem is encoded exclusively in two-body interactions through a single ancilla qubit, allowing periodic global spin-flip pulses to average out the longitudinal noise without disturbing the encoded ground state. Analytical arguments and numerical simulations on industrially motivated benchmarks demonstrate that modest pulse rates achievable on current hardware suffice to recover noise-free fidelity, and reveal a universal scaling of the final fidelity in terms of a generalized parameter combining noise amplitude and pulse interval.

On the algorithmic side, we investigate how the structure of the annealing Hamiltonian affects the minimum spectral gap that controls the annealing runtime. On the cost side, we show that retaining the native polynomial unconstrained binary optimization (PUBO) structure of combinatorial problems, rather than reducing them to the standard quadratic (QUBO) form, can significantly enlarge the minimum gap and reduce qubit requirements, with numerical studies on the paradigmatic 3-SAT problem revealing an exponential improvement in the gap-scaling exponent for certain problem instances. Finally, we study non-stoquastic annealing protocols with the ferromagnetic p -spin model and a geometrically local Ising model as benchmarks, and find that the regimes in which non-stoquasticity enlarges the spectral gap coincide with those of enhanced bipartite entanglement and non-stabilizerness of the instantaneous ground state, providing a link between annealing performance and the difficulty of classical simulation.

Taken together, these results support a clear path toward practical quantum annealing on near-term quantum hardware, combining a careful characterization of the physical platform, an active mitigation of the dominant noise channel, and a suitable design of the algorithm.

Contents

Introduction	1
List of publications	11
1 Trapped-ion quantum computing with magnetic gradient induced coupling	13
1.1 Ion traps	13
1.1.1 Overview of experimental ion traps	14
1.1.2 Ion crystal geometry and collective motion	16
1.1.3 Typical noise sources	19
1.2 Magnetic Gradient Induced Coupling (MAGIC)	20
1.2.1 Derivation of effective spin-spin coupling	20
1.2.2 Numerical example with $^{171}\text{Yb}^+$ ions	24
1.3 Implementing discrete quantum gates with trapped ions	26
1.3.1 Single-qubit gates	26
1.3.2 Two-qubit entangling gates	27
1.3.3 Implementing gates with MAGIC	30
2 Higher-order effects in MAGIC-based trapped-ion platforms	33
2.1 Higher-order contributions from the potential energy	34
2.1.1 General third-order expansion of the potential	34
2.1.2 Higher-order effects from the Coulomb repulsion	36
2.1.3 Influence of transversal phonon modes	41
2.1.4 Cubic anharmonicities of the external trapping potential	43
2.2 Contributions from magnetic field curvature	44
2.3 Comparison to existing schemes for generating higher-order interactions on trapped-ion platforms	47
3 Quantum optimization	49
3.1 Quantum annealing	49
3.1.1 From classical cost functions to cost Hamiltonians	50
3.1.2 Quantum annealing protocols	51
3.1.3 Adiabatic theorem	52
3.1.4 Characteristic energy scale	53
3.1.5 Optimal driving Hamiltonian strength	55
3.2 Benchmark models and combinatorial optimization problems	56
3.2.1 p -spin model	57
3.2.2 Geometrically local Ising model	58

3.2.3	Multiple Object Tracking	61
3.2.4	Cutting stock problem	65
3.2.5	3-SAT	69
3.2.6	Selected problems with natural PUBO formulations	72
3.3	Related quantum algorithms for solving optimization problems	73
3.3.1	Trotterized quantum annealing	73
3.3.2	Quantum approximate optimization algorithm (QAOA)	75
4	Noise mitigation for quantum annealing via dynamical decoupling	77
4.1	Effect of static disorder on quantum annealing performance	78
4.1.1	Fluctuations in the two-body couplings	78
4.1.2	Local field fluctuations	80
4.2	Applying dynamical decoupling pulses during annealing	81
4.2.1	Effect of dynamical decoupling on quantum annealing protocols	82
4.2.2	Analytical treatment via Magnus expansion	83
4.2.3	Modulating coupling strengths with dynamical decoupling	85
4.3	Numerical simulations of dynamical decoupling protocol	86
4.3.1	Energy scales and unit conventions	87
4.3.2	Generating time-dependent local-field noise	87
4.3.3	Numerical results for MOT, cutting stock and SK model	89
4.3.4	Universal scaling behavior	98
5	Improving quantum annealing performance	101
5.1	Utilizing higher-order interactions in the cost Hamiltonian	102
5.1.1	Reducing PUBO to QUBO	102
5.1.2	Energy scales and penalty strength	105
5.1.3	Structure of PUBO and QUBO formulations of 3-SAT	107
5.1.4	Scaling of the minimum energy gap	110
5.1.5	Implementing PUBO Hamiltonians on quantum hardware	112
5.1.6	Speedup of PUBO versus QUBO formulation	113
5.2	Inhomogeneous driving	114
5.3	Non-stoquastic quantum annealing	115
6	Quantum resources in non-stoquastic annealing	119
6.1	Entanglement and non-stabilizerness	120
6.1.1	Bipartite entanglement entropy	120
6.1.2	Stabilizer Rényi entropy	121
6.2	Numerical results for the p -spin model	122
6.2.1	Non-stoquastic annealing protocol for the p -spin model	122
6.2.2	Quantum resources across the annealing parameter space	123
6.2.3	Finite-size scaling along sinusoidal annealing paths	124
6.2.4	Entanglement and magic barriers	126
6.3	Numerical results for the geometrically local Ising model	127
6.3.1	Non-stoquastic annealing protocol	127
6.3.2	Transitions in the magnetization landscape	128
6.3.3	Microscopic description of the plateau structure	129
6.3.4	Quantum resources across the annealing parameter space	134
6.3.5	Finite-size scaling along annealing protocols	135
	Conclusions and Outlook	139

Bibliography	165
---------------------	------------

CONTENTS

Introduction

Quantum mechanics, developed in the early twentieth century, underpins a first generation of technologies including the laser, the transistor, and magnetic resonance imaging. While the operation of these devices is quantum mechanical, they do not generally directly exploit many-body quantum resources such as entanglement. Over the past three decades, a second generation has emerged in which such features are themselves treated as resources, with information encoded in quantum degrees of freedom [1–3]. These developments define the field of quantum technologies, conventionally organized into three closely related pillars, namely quantum sensing, quantum communication, and quantum computing [2]. Quantum sensing exploits the sensitivity of coherent quantum systems to external perturbations to enhance the precision of measurements, with realizations ranging from optical atomic clocks to magnetometers based on superconducting devices and nitrogen-vacancy centers [4]. Quantum communication uses nonclassical correlations to enable cryptography applications such as quantum key distribution, and more broadly underpins the development of quantum networks and a future quantum internet [5, 6]. The third pillar, quantum computing, forms the main subject of this thesis.

Foundations of quantum computing

The underlying idea of quantum computing is that a controlled quantum register of N qubits can efficiently implement dynamics whose classical simulation is believed to require resources exponential in N , opening the possibility of speedups for selected computational tasks [7]. This premise has motivated extensive research across academic, industrial, and government programs worldwide. The rigorous case for quantum advantage is, however, narrower than this activity might suggest: provable superpolynomial speedups over the best known classical algorithms are established only for a small number of problems [8–10], and current hardware remains far from the regime required for large-scale fault-tolerant computation [11, 12]. Useful applications therefore depend both on algorithm design matched to the strengths and limitations of available hardware and on continued improvements of the hardware itself.

The model of quantum computation that underpins the modern field has its origins in the early 1980s, with the observation that a controlled quantum system might be used to simulate other quantum systems efficiently [13, 14], in contrast to the apparent exponential cost of doing so on a classical computer. This idea was clarified by Deutsch [15], who introduced the universal quantum Turing machine, with the equivalent and now standard quantum circuit model formulated shortly thereafter [16, 17]. A series of contrived problems with black-box oracle separations between quantum and classical computation [18–20] culminated in Shor’s polynomial-time algorithms for integer factorization and discrete logarithms [21] and in Grover’s algorithm for unstructured search [22]. In parallel, the discovery of quantum error-correcting codes [23–25] and of fault-tolerant gate constructions [26], followed by the

threshold theorems [27–29], established that arbitrarily long quantum computations can in principle be carried out reliably given sufficiently low physical error rates.

In the circuit model, the basic computational unit is the qubit, a two-level quantum system whose pure state is a unit vector $\alpha|0\rangle + \beta|1\rangle$ in $\mathbb{C}^2$, with $\alpha, \beta \in \mathbb{C}$ and $|\alpha|^2 + |\beta|^2 = 1$. A register of N qubits has a state space of dimension 2^N and admits correlations, i.e. entanglement, that have no classical analogue. Computation is implemented by applying a sequence of unitary operations, drawn from a universal gate set, followed by projective measurements of the register. Two further computational models, formally distinct from the circuit model just described, have been developed: the adiabatic, or analog, model, in which the answer to a computational problem is encoded in the ground state of a final Hamiltonian reached by slow evolution from a simple initial Hamiltonian [30]; and the measurement-based, or one-way, model, in which computation is driven by a sequence of measurements on a fixed entangled resource state [31]. All three models are polynomially equivalent in their computational power [31, 32], but they differ substantially in the physical operations they require and, accordingly, in the hardware platforms best suited to implement them. The quantum annealing protocols studied in this thesis belong to the adiabatic family. A small number of algorithmic primitives recur across many known quantum algorithms: the quantum Fourier transform, together with its application in phase estimation, underlies a broad class of algorithms [33]; amplitude amplification provides the basic mechanism behind the known quadratic speedups [34]; and quantum teleportation [35] plays a central role both as a primitive in distributed quantum information processing and in fault-tolerant gate constructions.

Quantum algorithms and applications

The computational power of quantum computers is formalized by the complexity class **BQP**, the class of decision problems solvable by quantum computers in polynomial time with bounded error. Its proof-verification analogue is **QMA**, in which a polynomial-time quantum verifier checks a polynomial-size quantum witness. The canonical complete problem for **QMA** is the local Hamiltonian problem [36, 37]: given a Hamiltonian composed of a sum of few-body interaction terms, one must decide whether the ground-state energy lies below or above two promised thresholds separated by an inverse-polynomial gap.

The known evidence for quantum speedups is strongest in a few structured settings. The paradigmatic example is Shor’s algorithm, which solves factoring and discrete logarithms in polynomial time and more generally fits into the abelian hidden subgroup framework, built on the quantum Fourier transform and phase estimation [21, 38]. Grover’s algorithm occupies a different regime: built on amplitude amplification, it provides a generic quadratic speedup available for unstructured search, and this improvement is the optimal lower bound [22, 39]. Thus the established picture is not that quantum computers generically accelerate all hard computational tasks, but rather that they can exploit particular kinds of structure, while offering only limited improvements in structureless settings. Recent results show that the list of new problems with potential quantum speedups keeps expanding [40, 41]. Their current significance is mainly conceptual: they extend the range of problems for which superpolynomial quantum speedups are known or conjectured.

Among the applications of quantum computing, the most established theoretical case is in cryptography. A sufficiently large fault-tolerant quantum computer running Shor’s algorithm would compromise widely used public-key encryption based on factoring and discrete logarithms, which is one of the main motivations for post-quantum cryptography [42, 43]. Digital quantum simulation provides the cleanest constructive application case. It was the

original motivation for quantum computation, and is supported by a developed algorithmic toolbox [14, 44–46]. Target systems include the Fermi–Hubbard model, lattice gauge theories, quantum chemistry and non-equilibrium quantum dynamics [47, 48]. Recent experiments have begun to move this program from small demonstrations toward large-scale digital simulations. IBM’s 127-qubit kicked-Ising experiment provided an early example of pre-fault-tolerant many-body simulation with error mitigation, while also illustrating the direct competition with classical benchmarks [49, 50]. More directly application-oriented fermionic simulations include Phasecraft’s Fermi–Hubbard experiments on Google’s Willow processor and on Quantinuum’s trapped-ion H2 system, together with related IBM Heron experiments on spin excitation spectra and real-time Fermi–Hubbard dynamics, and Quantinuum’s simulation of superconducting pairing correlations on the Helios system [51–55]. These results do not by themselves establish practical quantum advantage, but they show that quantum simulation is becoming increasingly competitive with classical simulation methods.

Quantum machine learning [56–60] and optimization are more tentative application areas. Combinatorial optimization problems can often be encoded in Ising-type cost functions or more general local Hamiltonians, making them natural targets for analog quantum annealing and for gate-based variational algorithms such as QAOA [61–65]. These approaches are attractive because they match the native structure of many near-term devices, but there is no general evidence that they outperform the best classical methods on practically relevant problems. Benchmarking studies of annealers and analyses of QAOA illustrate the central difficulty: performance depends sensitively on instance structure, circuit depth, parameter optimization, and noise [66, 67]. Quantum machine learning faces analogous caveats. Algorithms for linear algebra subroutines such as HHL [68], and applications such as the quantum support vector machine [69], suggest exponential improvements; however, these speedups typically rely on assumptions about state preparation, data access and the efficient extraction of useful classical information [70]. In addition, dequantization results and trainability barriers such as barren plateaus show that many apparent advantages require careful qualification [71–73].

A conceptually distinct application of quantum computing is the sampling of equilibrium ensembles of classical statistical-mechanics models. While general near-quadratic speedups for Monte Carlo expectation estimation are known [74, 75], the potential quantum benefit here is qualitatively different: recasting sampling as energy minimization of a suitable Hamiltonian, quantum or classical, has been shown to give access to configurations that are practically unreachable by real-space Monte Carlo, due to topological constraints, as demonstrated for dense polymer melts [76–78] and for conformational transitions in biomolecules [79, 80].

Another type of quantum application which also does not rely on any algorithmic speedup is the generation of certified random numbers. In certified-randomness protocols, the inability of classical computers to efficiently reproduce the output distribution of random quantum circuits allows a classical client to verify that fresh, unpredictable bits have been generated by a remote untrusted quantum device, a task provably impossible classically [81] and recently demonstrated on a trapped-ion processor [82].

Quantum advantage

Beyond constructive applications, a complementary line of research has aimed at the question of identifying the most accessible computational tasks for which a near-term quantum device can be shown to outperform any classical alternative. An idealised demonstration would jointly satisfy three criteria [83–85]: feasible implementation on current hardware,

average-case classical hardness, and an efficient verification procedure, ideally but not necessarily tied to a useful task. Existing experiments meet at most two of these requirements simultaneously, and progress in classical algorithms repeatedly redraws the boundary between accessible and inaccessible regimes.

The experiments in this direction have mainly focused on sampling-based protocols. The first claim of quantum computational supremacy was reported on Google’s 53-qubit Sycamore processor, which sampled from random quantum circuits and was validated through the linear cross-entropy benchmark [86], building on average-case hardness conjectures developed earlier for boson and random circuit sampling [87–90]. Photonic platforms followed with boson sampling demonstrations [91–93], and superconducting random-circuit sampling was extended [94, 95]. Subsequent generations have pushed deeper into the regime beyond the reach of classical methods. A qualitatively different demonstration was recently reported on Willow using a second-order out-of-time-order correlator (OTOC) protocol [96]: rather than a sampling distribution, the experiment outputs an expectation value that is reproducible and can be checked on independent quantum hardware, providing the first claim of a quantum computational advantage on a verifiable task.

Classical simulation methods

Quantum advantage claims have to be measured against steadily improving classical simulation methods, which give a bound on the regime of genuine quantum advantage and supply the analysis tools for quantum protocols. Tensor network methods [97–100] exploit area-law entanglement structure: the bond dimension of a matrix product state (MPS) required for a faithful quantum state representation grows with the amount of entanglement in the system, which constitutes the central bottleneck of the approach. Targeted tensor network simulations have repeatedly closed apparent gaps between quantum experiments and the best classical alternatives, most prominently for IBM’s 127-qubit kicked Ising experiment [49], where the reported observables were reproduced more accurately and at modest computational cost [50, 101, 102]. A second class of methods is rooted in the stabilizer formalism: Clifford circuits (composed solely of Hadamard, CNOT and phase gates) are efficiently classically simulable independent of the amount of entanglement [103, 104], and this can be extended to circuits with a controlled budget of non-Clifford operations, with simulation cost scaling with the nonstabilizerness of the final state [105, 106]. Pauli-path or Pauli propagation methods estimate expectation values of local observables in noisy circuits [102, 107–109]. Finally, matchgate or fermionic Gaussian circuits, equivalent under the Jordan–Wigner transformation to noninteracting fermion dynamics, form a separate efficiently simulable class [110–112].

A unifying theme emerges across these methods, namely that the cost of classical simulation is controlled by a small set of nonclassical resources. Entanglement entropy bounds tensor network approaches [113], nonstabilizerness or magic limits stabilizer-based approaches [114, 115], and the analogous fermionic resource bounds matchgate-based approaches [116]. These quantities thereby serve as quantitative diagnostics of where on the boundary between classical and quantum computation a given protocol resides, and they have correspondingly become objects of intrinsic interest in quantum many-body and quantum information theory. Bipartite entanglement entropy and stabilizer Rényi entropies [117] along non-stoquastic annealing paths return as central observables later in this thesis.

Quantum error correction and noise mitigation

The above-described quantum resources are not the only important factor when it comes to demonstrating quantum advantage for useful computations. The hardware available to realize the protocols just discussed is itself imperfect, and the present research is largely focused on finding strategies to contend with this fact. The currently accessible quantum computing regime was named the noisy intermediate-scale quantum (NISQ) era by Preskill [11], denoting devices with on the order of 10^1 – 10^2 noisy qubits, large enough that brute-force classical simulation is not always feasible but short of the resources required to support quantum error correction at scale. In response to incremental experimental progress, the community has more recently begun to distinguish an early fault-tolerant phase [118] and the long-term target of fault-tolerant application-scale quantum (FASQ) devices [9], in which encoded logical qubits support computations of practical interest.

Three complementary strategies for managing noise have emerged, and they differ qualitatively in their resource requirements and effects. Quantum error correction (QEC) redundantly encodes logical information in many physical qubits and uses repeated syndrome measurements together with a classical decoder to detect and correct errors; its guarantee, the threshold theorem [27–29], is that arbitrarily long computations can be carried out reliably provided the physical error rate is below a code- and architecture-dependent threshold value, at the cost of substantial qubit overhead. Quantum error mitigation (QEM) instead leaves the hardware uncorrected and trades sampling overhead for accuracy, by combining many runs of noisy circuits into improved estimates of expectation values [119]. QEM methods are in the literature often distinguished from quantum error suppression, which operates at the level of control pulses, using tailored coherent operations to effectively decouple the system from its environment without measurement or post-processing. In this thesis we use the terms noise mitigation and suppression interchangeably.

Mitigation methods have matured in parallel with NISQ hardware. Zero-noise extrapolation [120–122] estimates the noiseless expectation value by deliberately amplifying the noise rate in a controlled fashion and then extrapolating back to zero; probabilistic error cancellation [121, 123] attempts to invert a sufficiently well-characterized noise channel, yielding an unbiased estimator at the price of a sampling overhead that grows exponentially with circuit volume. Virtual distillation [124, 125] suppresses incoherent errors by estimating expectation values using multiple noisy state copies. Learning-based approaches such as Clifford data regression [126] train a correction map on classically simulable instances, and symmetry verification [127, 128] discards measurement outcomes that violate conserved quantities of the ideal evolution. The IBM kicked-Ising experiment discussed earlier [49] relied centrally on these techniques. Fundamental limits constrain the strategy: the sampling overhead of any unbiased mitigation protocol grows at least exponentially with the circuit volume [129, 130], meaning that QEM extends the useful regime of NISQ hardware but does not provide a scalable path to fault tolerance.

Dynamical decoupling occupies a distinct position. Rather than correcting or post-processing errors, it suppresses the system-environment coupling at the Hamiltonian level by applying carefully tailored sequences of control pulses, so that the time-averaged interaction with the environment is effectively cancelled to a chosen order. The technique originates in nuclear magnetic resonance [131]. Its formulation in the language of quantum information [132, 133] established that a finite number of pulses can average a generic system-bath interaction to zero, and a substantial body of subsequent work has produced higher-order and more robust pulse sequences [134–136]. Dynamical decoupling is now a routine protocol on essentially every major hardware platform and has been used to extend qubit coher-

ence times by orders of magnitude on trapped-ion [137, 138], superconducting [139, 140], nitrogen-vacancy centre [141], and neutral-atom systems [142]. The scheme has also been extended to analog algorithms, where the relevant control operations commute non-trivially with a time-dependent Hamiltonian [143, 144].

The long-term goal of the field is full quantum error correction. The underlying idea is to encode each logical qubit redundantly across many physical qubits, in such a way that low-weight physical errors (acting on a small number of physical qubits at the same time) are detected by repeated parity-type (syndrome) measurements and reversed before they can corrupt the encoded quantum information. The relevant resilience parameter is the code distance d , the minimum weight of an undetectable error and intuitively a measure of how much protection the code provides. Once the physical error rate falls below a code-dependent threshold value, beyond which adding more physical qubits actually helps rather than hurts, increasing d suppresses the logical error rate exponentially, and a universal fault-tolerant gate set acting directly on logical qubits in principle supports arbitrarily long computations. Translating this picture into hardware has been the object of decades of work, and a number of recent experiments mark qualitative progress: surface-code demonstrations on Google’s Willow processor crossed the threshold for the first time, with the logical error rate suppressed by a roughly constant factor at each increase of the code distance [145]; Trapped-ion experiments at Quantinuum have demonstrated repeatable error correction on a $[[7, 1, 3]]$ Steane code [146]; reconfigurable neutral-atom arrays at Harvard and QuEra have implemented logical processors with hundreds of physical qubits [147, 148]. Low-density parity-check (qLDPC) constructions, including the bivariate bicycle codes proposed for superconducting hardware [149], suggest a path toward substantially lower qubit overheads for FASQ devices. None of these results is itself sufficient for application-scale fault tolerance, but together they shift the focus from whether the error threshold can be crossed to how fast the required resources for FASQ devices can be reduced. The characteristics of a physical platform, e.g. noise, connectivity, and control, continue to shape both the algorithms that are realistic to implement and the mitigation and correction strategies that bridge the gap to full fault-tolerance.

Physical realizations

The algorithms and error mitigation/correction strategies surveyed above are realized on a heterogeneous landscape of different hardware platforms, each with characteristic strengths and weaknesses regarding qubit count, connectivity, gate fidelity, coherence time, and suitability for scalable error correction. The remainder of this section briefly surveys the leading platforms. As many of the protocols discussed in this thesis are tailored to the trapped-ion platform with MAGIC, we defer a more detailed treatment of this platform to Chapters 1 and 2.

Superconducting circuits are at present one of the most mature gate-based platforms by qubit count. The dominant qubit design is the transmon [150], an anharmonic oscillator formed by a Josephson junction shunted by a large capacitance. The fluxonium [151] provides a higher-anharmonicity alternative that has recently received renewed attention, in which the Josephson junction is shunted by an inductor instead. Two-qubit gate fidelities on fixed-frequency or tunable-coupler transmon devices have reached or even surpassed 99.9% across multiple groups [152–154], and recent fluxonium demonstrations have reached comparable performance [155]. IBM has scaled transmon arrays through the 127-qubit Eagle, the 133-qubit Heron processor introducing tunable couplers, and the 1,121-qubit Condor chip [156]. Google has demonstrated the first below-threshold operation of the surface code on its

105-qubit superconducting Willow processor, as discussed above [145]. Other companies developing superconducting quantum computers include Rigetti, IQM and Alice & Bob. A separate line of superconducting hardware uses flux qubits to realize transverse-field Ising Hamiltonians for analog quantum annealing, as implemented by D-Wave systems.

Trapped ions are the second leading gate-based platform discussed here and the reference platform considered in this thesis [157]. Quantum information is encoded in long-lived hyperfine or optical levels of individual ions, addressed by laser or microwave fields, and entangled through collective vibrational modes of the ion crystal which allow long-range all-to-all coupling between ions. The standard entangling gate schemes are the Cirac-Zoller [158] and Mølmer-Sørensen [159, 160] gates. A promising variant on the trapped-ion platform is the Magnetic Gradient Induced Coupling (MAGIC) scheme, which is solely controlled by microwaves for single-qubit rotations and uses an external magnetic field gradient to couple ions [161–163]. Two-qubit gate fidelities first crossed 99.9% in laser-driven hyperfine qubits in $^{43}\text{Ca}^+$ [164, 165], with microwave-controlled schemes [166, 167] reaching comparable performance, and most recent results from electronically-controlled ion traps reporting two-qubit gate errors below 10^{-4} [168]. Trapped-ion architectures range from long one-dimensional ion chains, to QCCD ion-shuttling approaches in the Quantinuum processors, two-dimensional Penning-trap crystals, and proposed photonic interconnects between modules. Commercial developers include Quantinuum, IonQ, eleQtron, and AQT, alongside long-running academic and government laboratory programs.

Neutral atoms in optical tweezer arrays have advanced rapidly from analog quantum simulation toward programmable, gate-based, and early fault-tolerant operation. Qubits are encoded in hyperfine or nuclear-spin states of trapped atoms, with two-qubit gates mediated by the strong dipole-dipole interactions of highly-excited Rydberg states [169, 170]. The reconfigurability of the array supports logical-qubit architectures with effectively non-local connectivity, a feature exploited in several of the recent fault-tolerance demonstrations. Representative milestones include QuEra’s 256-atom Aquila analog processor; Atom Computing’s array of over 1000 ^{171}Yb nuclear-spin qubits [171]; the Harvard-MIT-QuEra programmable logical processor operating up to 48 logical qubits on 280 physical atoms [147], and a follow-up implementation of key components of a fault-tolerant universal architecture using up to 448 atoms [148]. Atom Computing in collaboration with Microsoft has separately reported entanglement of 24 logical qubits and execution of the Bernstein-Vazirani algorithm on 28 logical qubits with active error correction [172].

Photonic platforms operate at room temperature, with coherence times limited primarily by photon loss, and are naturally compatible with quantum networking. The central obstacle is the absence of deterministic two-photon entangling gates: standard linear-optical schemes implement them only probabilistically [173], which has motivated fault-tolerant architectures based on cluster and fusion-network resource states [174]. Xanadu has used programmable squeezed-light interferometers in boson sampling experiments [93], and Psi-Quantum recently benchmarked a manufacturable silicon-photonic platform with integrated single-photon sources and detectors approaching fault-tolerant thresholds [175].

Several further platforms remain at earlier stages of development. Semiconductor spin qubits in silicon and germanium quantum dots are pursued for their compatibility with established CMOS fabrication [176] with several demonstrations of single- and two-qubit gates above the surface-code threshold [177, 178]. Nitrogen-vacancy centres in diamond combine room-temperature operation with applications in quantum sensing [179]. Microsoft has long pursued a distinct route based on topologically protected qubits encoded in Majorana zero modes; in February 2025, the company announced the Majorana 1 device, accompanied by a publication reporting an interferometric fermion-parity measurement in nanowire

devices [180]. The peer-reviewed publication itself does not demonstrate Majorana zero modes in the reported devices, as the journal’s editorial team and the peer-review explicitly note, and the broader claims have been the subject of substantial subsequent scientific scrutiny [181].

All of these platforms differ substantially in their underlying physics and in the resulting trade-offs between connectivity, gate fidelity, gate speed, coherence time, and scalability. This heterogeneity motivates the platform-specific perspective adopted in this thesis: rather than studying quantum annealing protocols in the abstract, the focus is on algorithmic and noise-mitigation design tailored to a specific trapped-ion realization, namely the MAGIC platform.

Quantum annealing

The quantum algorithm of central interest in this thesis is quantum annealing [30, 64, 182, 183], an analog approach to combinatorial optimization rooted in the model of adiabatic quantum computation, with the terms often being used interchangeably. Building on classical simulated annealing [184], the underlying idea is to encode the solution of a classical optimization problem one wishes to solve into the ground state of a final cost Hamiltonian $\hat{H}_{\text{cost}}$ and to reach it by slow interpolation from a simple driver $\hat{H}_{\text{drive}}$ whose ground state is known and easy to prepare. The standard annealing schedule reads

$$\hat{H}(s) = (1 - s) \hat{H}_{\text{drive}} + s \hat{H}_{\text{cost}}, \quad s \in [0, 1], \quad (1)$$

with s a dimensionless time parameter. The adiabatic theorem [185–188] guarantees that, for a sufficiently slow evolution, the system follows the instantaneous ground state of $\hat{H}(s)$, with the required runtime scaling as the inverse squared minimum spectral gap along the evolution. The central question for the efficiency of any annealing protocol is therefore how the minimum gap scales with problem size. Adiabatic quantum computation in its general form was shown to be polynomially equivalent to the circuit model [32], although hardware annealers such as those developed by D-Wave Systems realize only Hamiltonians of a restricted transverse-field Ising form and are therefore not universal in this sense [64, 183]. The interest of quantum annealing thus rests largely on heuristic and experimental grounds, and on the convenient compatibility of its native problem structure with classical combinatorial optimization, which is typically cast in quadratic (QUBO) or higher-order polynomial (PUBO) unconstrained binary form [189].

A central distinction within the family of annealing protocols is between stoquastic and non-stoquastic forms [190]. The standard transverse-field Ising annealer is stoquastic in the computational basis, meaning that all of its off-diagonal matrix elements are real and non-positive, making it sign-problem-free for Monte-Carlo simulations. Including non-stoquastic terms, that is, off-diagonal matrix elements of positive real part such as antiferromagnetic transverse XX couplings, yields a potentially richer class of annealing protocols whose simulation generically suffers from a sign problem and can recover **QMA**-completeness of the corresponding local Hamiltonian problem [191].

The principal obstruction to standard transverse-field annealing is closing of the spectral gap at first-order quantum phase transitions, which produces an exponentially small minimum gap and forces an exponentially large runtime [192, 193]. Together with restricted hardware connectivity, finite-temperature decoherence, and the absence of any provable speedup over the best classical algorithms [66], these limitations have driven extensive work on improvements to the standard scheme, two directions of which are mainly studied in this thesis. The first enriches the cost Hamiltonian by retaining native k -body interactions

with $k > 2$, which arise naturally for many combinatorial problems and would otherwise be quadratized at the cost of ancillary slack variables and large penalty terms [189, 194]. The second modifies the driving Hamiltonian by introducing non-stoquastic catalysts during an intermediate window of the schedule. Seki and Nishimori showed analytically that such catalysts convert the first-order transition of the p -spin model into a second-order one and replace the exponentially closing gap by a polynomially closing one [195, 196]. The picture is, however, strongly problem-dependent: numerical studies on random spin glasses observed non-stoquastic improvements only on a small fraction of instances, with the gain attributed in part to diabatic transitions rather than to a uniform enlargement of the gap [197].

Quantum annealing protocols have been implemented across different platforms. The most mature realization is the family of superconducting flux-qubit annealers developed by D-Wave Systems, which has reached several thousand qubits [198]. Reconfigurable Rydberg-atom arrays have implemented analog annealing protocols for the Maximum Independent Set problem on graphs of hundreds of vertices [199], and annealing schedules have been digitized via Trotterization on superconducting processors [200]. Trapped ions themselves are an established platform for analog quantum simulation of spin models [201, 202], with full programmability of the spin-spin couplings and natural compatibility with the structure of annealing Hamiltonians. The microwave-controlled MAGIC architecture introduced earlier is especially promising for quantum optimization, as it offers all-to-all always-on Ising-like interactions between all qubits and constitutes the hardware setting we consider in this thesis.

Scope and structure of this thesis

The work presented in this thesis aims to bring quantum annealing closer to practical application on near-term trapped-ion hardware, and is organized along two complementary lines of investigation. The first line is hardware-aware and tailored to the MAGIC trapped-ion platform. We perform a perturbative analysis of higher-order corrections to the MAGIC spin-spin couplings, including Coulomb anharmonicities, magnetic-field curvature, and axial-transversal phonon couplings [203], identifying the dominant terms that need to be accounted for when scaling to larger ion chains. We further design and analyze a dynamical decoupling protocol tailored to analog quantum annealing on this platform, in which an ancilla encoding maps local fields onto two-body interactions so that periodic global spin-flip pulses suppress the dominant noise terms without affecting the encoded problem [144]. The second line is algorithmic and mostly hardware-agnostic, focusing on how the structure of the annealing Hamiltonian affects performance through the minimum spectral gap. On the side of the cost Hamiltonian, we show on 3-SAT benchmarks that retaining the native polynomial unconstrained binary optimization (PUBO) structure can improve the gap-scaling exponent relative to the standard quadratic (QUBO) reduction [203]. On the side of the driver, we study annealing protocols supplemented by non-stoquastic catalysts on the paradigmatic p -spin and a geometrically local Ising model, and find that the regimes in which non-stoquasticity enlarges the spectral gap coincide with those of enhanced bipartite entanglement entropy and non-stabilizerness along the annealing path [204], linking the complexity-theoretic motivation for non-stoquasticity to a resource-theoretic notion of being hard to simulate on classical computers.

The remainder of this thesis is organized as follows. Chapter 1 introduces the trapped-ion background underlying the hardware-related part of the work, covering experimental ion-trap architectures, the equilibrium positions and phonon modes of linear ion chains, the dominant noise sources for magnetically sensitive qubits, the derivation of the effective spin-

spin coupling in the MAGIC scheme, and the implementation of single- and two-qubit gates with trapped ions. Chapter 2 discusses higher-order corrections to the MAGIC couplings arising from Coulomb anharmonicities and magnetic-field curvature, and identifies a phonon-occupation-dependent longitudinal field as the most relevant residual noise term emerging from this analysis. Chapter 3 discusses the theoretical framework of quantum annealing, covering the encoding of optimization problems into cost Hamiltonians, the adiabatic theorem, and the role of the relevant energy scales. It then introduces the combinatorial benchmarks used in the rest of the thesis, including the ferromagnetic p -spin model, a geometrically local Ising model, industrially motivated problems, and the canonical 3-SAT problem, and it closes with a brief review of Trotterized quantum annealing and QAOA as digital counterparts of the analog protocols. Chapter 4 develops a dynamical decoupling protocol tailored to analog quantum annealing on the MAGIC platform, based on periodic global spin-flip pulses combined with an ancilla encoding that maps local fields onto two-body interactions, and shows through a Magnus expansion and numerical simulations on industrially relevant benchmarks that modest pulse rates suffice to recover noise-free fidelity under realistic noise. Chapter 5 addresses the design of the cost Hamiltonian itself, comparing the direct PUBO formulation of 3-SAT with the standard QUBO reduction and showing that retaining the native higher-order structure improves the gap-scaling exponent. It further discusses inhomogeneous driving and non-stoquastic catalysts as complementary strategies to enhance annealing performance. Chapter 6 investigates the connection between non-stoquastic catalysts and the quantum resources generated along the annealing path. After reviewing bipartite entanglement entropy and stabilizer Rényi entropy together with their efficient numerical evaluation in the presence of symmetries, we analyze the ferromagnetic p -spin model and a geometrically local Ising model, and find that the regimes in which non-stoquasticity improves the spectral gap coincide with those of enhanced entanglement and non-stabilizerness. Finally, the last part summarizes the main results of this thesis and outlines directions for future research.

List of publications

Publications included in this thesis

Chiara Capecci, **Sebastian Nagies**, Naga Dileep Varikuti, Philipp Hauke, *Quantum resources in non-stoquastic annealing*. arXiv (2026), doi:10.48550/arXiv.2606.09995

Discussed in Chapters 3, 5 and 6. Contribution: Shared first author with C. Capecci; numerical calculations together with C. Capecci; conception and writing together with co-authors

Sebastian Nagies, Chiara Capecci, Marcel Seelbach Benkner, Javed Akram, Sebastian Rubbert, Dimitrios Bantounas, Michael Moeller, Michael Johanning, Philipp Hauke, *Practical Noise Mitigation for Quantum Annealing via Dynamical Decoupling: Toward Industry-Relevant Optimization using Trapped Ions*. Phys. Rev. A 113, 012621 (2026), doi:10.1103/mfqh-jtrl

Discussed in Chapters 3 and 4. Contribution: Analytical calculations together with C. Capecci; numerical calculations together with C. Capecci and M. Seelbach Benkner; conception and writing together with co-authors

Sebastian Nagies, Kevin T. Geier, Javed Akram, Dimitrios Bantounas, Michael Johanning, Philipp Hauke, *Boosting quantum annealing performance through direct polynomial unconstrained binary optimization*. Quantum Sci. Technol. 10 035008 (2025), doi:10.1088/2058-9565/adcae6

Discussed in Chapters 3 and 5. Contribution: Numerical calculations; analytical calculations together with K. T. Geier; conception and writing together with co-authors

Sebastian Nagies, Kevin T. Geier, Javed Akram, Junichi Okamoto, Dimitrios Bantounas, Christof Wunderlich, Michael Johanning, Philipp Hauke, *The role of higher-order terms in trapped-ion quantum computing with magnetic gradient induced coupling*. Quantum Sci. Technol. 10 025051 (2025), doi:10.1088/2058-9565/adc1fe

Discussed in Chapters 1 and 2. Contribution: Analytical and numerical calculations; conception and writing together with co-authors

Other publications

Emiliano Tolotti, **Sebastian Nagies**, Davide Pastorello, Enrico Blanzieri, *Making Quantum Neural Networks Practical: A Study of Measurement Efficiency*. in preparation

Sebastian Nagies, Emiliano Tolotti, Davide Pastorello, Enrico Blanzieri, *Enhancing Expressivity of Quantum Neural Networks Based on the SWAP test*. arXiv (2025), doi:10.48550/arXiv.2506.16938

Sebastian Nagies, Botao Wang, Adam C. Knapp, André Eckardt, Nathan L. Harshman, *Beyond braid statistics: Constructing a lattice model for anyons with exchange statistics intrinsic to one dimension*. SciPost Phys. 16, 086 (2024), doi:10.21468/SciPostPhys.16.3.086

Chapter 1

Trapped-ion quantum computing with magnetic gradient induced coupling

Trapped ions are among the most mature platforms for digital quantum computation, combining long coherence times with high single- and two-qubit gate fidelities. Within this broad family of architectures, the *magnetic gradient induced coupling* (MAGIC) scheme stands out for its use of microwave radiation driving fields in place of optical lasers, exploiting a static magnetic field gradient both for individual addressing and for generating an always-on Ising interaction between the ions. This native $\sigma_Z\sigma_Z$ coupling makes MAGIC particularly well suited as a hardware reference for quantum annealing, which provides the broader context of this thesis. Throughout this work, MAGIC therefore serves as our concrete platform of choice: the higher-order corrections analyzed in Chapter 2 and the noise mitigation strategies developed in Chapter 4 are formulated specifically with this platform in mind, and the experimental parameters used throughout correspond to the MAGIC implementation pursued by eleQtron.

The present chapter introduces the background needed for these later developments. Section 1.1 reviews experimental ion-trap architectures, the equilibrium structure and collective motion (phonon modes) of linear ion chains, and the dominant noise sources relevant to magnetically sensitive qubits. Section 1.2 derives the effective spin-spin coupling generated in the MAGIC setup by the magnetic field gradient and provides numerical estimates of the achievable coupling strengths. Section 1.3 describes how single- and two-qubit gates are realized on this platform as well as on laser-based platforms in general.

The discussion of the MAGIC platform, including the derivation of the spin-spin coupling, as well as the discussion of gate implementations for that platform, is based on material published in Ref. [205].

1.1 Ion traps

This section provides the experimental and theoretical background on ion traps needed for the rest of the chapter. Section 1.1.1 gives a brief overview of the two main classes of trap architectures, Penning and Paul traps, and their respective roles in quantum simulation and digital quantum computing. Section 1.1.2 discusses the equilibrium configuration of a one-

dimensional ion chain and its collective phonon modes, which mediate the effective spin-spin coupling derived in Sec. 1.2.1. Section 1.1.3 summarizes the dominant noise sources affecting trapped-ion qubits, with particular attention to the magnetic-field fluctuations relevant for the magnetically sensitive states used in the MAGIC scheme.

1.1.1 Overview of experimental ion traps

Penning traps

A Penning trap confines ions (or more generally charged particles) using a uniform static magnetic field $\mathbf{B} = B\hat{z}$ for radial confinement and a static quadrupole electric potential for axial confinement [206, 207]. Hyperbolically shaped electrodes generate the potential $\Phi(r, z) = \frac{V_0}{2d^2}(z^2 - r^2/2)$, where $d^2 = \frac{1}{2}(z_0^2 + r_0^2/2)$ is a characteristic trap dimension determined by the minimal axial distance to the electrodes z_0 and the minimal radial distance r_0 [206]. This creates an axial harmonic well which is not radially confining. The Lorentz force due to $\mathbf{B}$ overcomes this instability, bending ion trajectories into closed orbits.

Several factors have historically limited the use of Penning traps for digital quantum computing. Multi-ion crystals rotate due to the Lorentz force, making individual ion addressing extremely challenging [208]. The strong magnetic field ($B \sim 1\text{--}5$ T, requiring a superconducting solenoid) introduces large Zeeman splittings, resulting in qubit transitions of tens or hundreds of GHz [209], and magnetic-field fluctuations directly cause qubit dephasing. For these reasons, commercial digital trapped-ion quantum processors have so far followed Paul-trap architectures rather than conventional rotating-crystal Penning-trap architectures.

Penning traps have, however, proven powerful for quantum simulation. The NIST group has used 2D Coulomb crystals of $^9\text{Be}^+$ ions to engineer tunable Ising interactions across $\sim 200\text{--}300$ spins [201], demonstrate spin squeezing with up to 219 ions [210], and probe information scrambling via out-of-time-order correlators [211].

A recent development is the demonstration of a micro-fabricated surface-electrode Penning trap by the ETH Zürich TIQI group [212], building on their theoretical proposal for scalable 2D micro-Penning arrays [213]. A single $^9\text{Be}^+$ ion was trapped above a chip surface at $B = 3$ T, with full control of spin and motional degrees of freedom and arbitrary 2D transport with low motional heating. This approach enables junction-free 2D ion transport using only DC voltage adjustments, eliminates the rotating-crystal problem by trapping one ion per site, and dissipates zero RF power. The critical open milestones are demonstrating a two-qubit gate and operating a multi-site array with controllable inter-site coupling.

Paul traps

Paul traps confine ions through a combination of static and oscillating electric fields. Earnshaw’s theorem forbids a stable electrostatic minimum for a charged particle in free space, so a time-varying RF quadrupole field is used: at each instant in time the field is confining along one direction and anti-confining along the orthogonal one, but the rapid oscillation generates an effective time-averaged pseudopotential that confines the ion in all three dimensions [214]. The resulting ion dynamics separates into a slow harmonic secular motion at a frequency set by the RF amplitude, the RF drive frequency, and the trap geometry, superimposed on a fast small-amplitude driven oscillation at the RF frequency known as micromotion [215]. An ion sitting exactly on the RF null experiences no micromotion; the intrinsic micromotion of a cooled ion arises only because its secular motion periodically displaces it from the null. Stray DC fields, however, push the ion’s equilibrium position away

from the null and produce excess micromotion, which cannot be reduced by cooling and is a dominant source of gate errors, second-order Doppler shifts, and AC Stark shifts [215].

Macroscopic linear Paul traps. The most common geometry for trapped-ion quantum computing is the linear Paul trap, in which four parallel rod or blade electrodes provide radial RF confinement and segmented DC endcap electrodes generate axial confinement [216, 217]. The resulting RF null is a line along the trap axis, on which laser-cooled ions crystallize into a 1D Coulomb chain, for suitable trap parameters and ion numbers (see Sec. 1.1.2). Typical operating parameters are RF drive frequencies of 10–100 MHz, RF voltages of 100–1000 V, ion-electrode distances of 0.1–1 mm, and secular frequencies of 0.1–5 MHz. Macroscopic traps offer deep confining potentials, good optical access, and low motional heating rates, properties that enabled the first demonstrations of single- and two-qubit gate fidelities well above the fault-tolerance threshold [164, 218]. Scaling beyond a few tens of qubits in a single chain is difficult: as the chain lengthens, the transverse motional spectrum becomes dense and its lowest modes soften, eventually triggering a zigzag structural instability, while inter-ion spacings shrink [219] (see discussion in Sec. 1.1.2). These effects degrade individual addressing and gate fidelity.

Surface-electrode (chip) traps. A complementary approach proposed by Chiaverini et al. [220] and first demonstrated by Seidelin et al. [221] places all trap electrodes in a single plane on a chip substrate, trapping ions in vacuum at typical heights of 30–200 μm above the surface. The planar geometry yields shallower trap depths than macroscopic traps but is fully compatible with semiconductor microfabrication. This fabrication route naturally allows for segmented traps, in which the static control electrodes are divided into many independently controllable sections [222]. Similar segmentation can also be used in macroscopic Paul traps [223, 224], and more recently in microfabricated Penning-trap designs [212, 213], but in trapped-ion quantum computing the term most commonly refers to segmented Paul chip traps. By dynamically varying the segment voltages, one can shape the axial potential to shuttle, split, merge, and reorder ions, enabling complex two-dimensional layouts with junctions, storage regions, and dedicated gate zones. These operations provide the geometric and control substrate for practical multi-zone QCCD architectures [225], as well as for on-chip integration of photonic waveguides, microwave structures, and control electronics on multi-layer wafer processes [226].

State of the art and recent developments. Trapped-ion chip-scale platforms continue to advance along two complementary directions: integrated microwave/electrode-based control and large-scale laser-based QCCD architectures. The use of microwave fields and magnetic field gradients generated by trap-integrated electrodes was originally proposed by Mintert and Wunderlich [161] and demonstrated in Refs. [163, 227]. Recent experiments within this paradigm have achieved extremely low gate errors. Smith et al. [228] reported a minimum single-qubit Clifford error of $1.5(4) \times 10^{-7}$ on a $^{43}\text{Ca}^+$ hyperfine clock qubit, while Hughes et al. [168] demonstrated a two-qubit gate error of $8.4(7) \times 10^{-5}$ on $^{40}\text{Ca}^+$ without ground-state cooling. In parallel, system-level scaling has been pursued using laser-based architectures. Quantinuum’s Helios processor implements a multi-zone QCCD design with 98 $^{137}\text{Ba}^+$ qubits and a reported two-qubit gate fidelity of approximately 99.921% [222]. Taken together, these results demonstrate that chip-based trapped-ion systems now achieve state-of-the-art component performance while following distinct hardware approaches.

1.1.2 Ion crystal geometry and collective motion

When multiple ions are confined in a Paul trap, they repel each other via the Coulomb interaction and, if sufficiently cooled down, arrange themselves in a crystalline structure whose geometry depends on the external trapping frequencies [219, 229]. In this section, we discuss the equilibrium configuration of N identical ions of mass M and charge $Q = Ze$ in a three-dimensional harmonic trapping potential, with particular attention to the experimentally relevant regime where the ions form a one-dimensional chain along the weakly confining axial direction. We then describe the collective vibrational modes (phonon modes) of the crystal, which play a central role for calculating the effective spin-spin interactions for the MAGIC platform derived in Sec. 1.2.1.

Potential energy and equilibrium positions

We consider a harmonic trapping potential characterized by angular frequencies ω_x , ω_y in the two transverse directions and ω_z in the axial direction, with $\omega_x, \omega_y \gg \omega_z$. The position of the n -th ion is denoted by (x_{n1}, x_{n2}, x_{n3}) , where the third coordinate corresponds to the axial direction and the ions are numbered in order of increasing axial position. The total potential energy of the system comprises the confining harmonic potential and the pairwise Coulomb repulsion between ions [219, 229],

$$V = \frac{M}{2} \sum_{n=1}^N \sum_{i=1}^3 \omega_i^2 x_{ni}^2 + \frac{Q^2}{8\pi\epsilon_0} \sum_{\substack{n,m=1 \\ m \neq n}}^N \left[\sum_{i=1}^3 (x_{ni} - x_{mi})^2 \right]^{-1/2}, \quad (1.1)$$

where $\omega_1 = \omega_x$, $\omega_2 = \omega_y$, $\omega_3 = \omega_z$, and ϵ_0 is the vacuum permittivity.

For sufficiently strong transverse confinement, the ions are tightly bound in the x_1 and x_2 directions and equilibrate along the axial direction, forming a linear chain with transverse equilibrium positions $\bar{x}_{n1} = \bar{x}_{n2} = 0$. The axial equilibrium positions $\bar{x}_{n3} \equiv x_n^{(0)}$ are determined by the condition

$$\left. \frac{\partial V}{\partial x_{n3}} \right|_{\bar{x}} = 0, \quad n = 1, \dots, N. \quad (1.2)$$

It is convenient to introduce dimensionless coordinates. Defining the characteristic length scale

$$\ell = \left(\frac{Q^2}{4\pi\epsilon_0 M \omega_z^2} \right)^{1/3}, \quad (1.3)$$

and the dimensionless equilibrium positions $u_n = x_n^{(0)}/\ell$, the equilibrium conditions (1.2) reduce to a set of N coupled algebraic equations [219]

$$u_m - \sum_{n=1}^{m-1} \frac{1}{(u_m - u_n)^2} + \sum_{n=m+1}^N \frac{1}{(u_m - u_n)^2} = 0, \quad m = 1, 2, \dots, N. \quad (1.4)$$

While these equations can be solved analytically for $N = 2, 3$ (see Ref. [219]), they have to be solved numerically for larger N . The equilibrium positions are symmetric about the trap center ($u_n = -u_{N+1-n}$), and the minimum inter-ion spacing, which occurs at the center of the chain, can be approximated by [219]

$$u_{\min}(N) \approx \frac{2.018}{N^{0.559}}. \quad (1.5)$$

Normal modes of the collective oscillation

For ions at sufficiently low temperature, the displacement from equilibrium is small, and the potential energy can be expanded in a Taylor series around the equilibrium positions. Denoting the axial displacement of the n -th ion from equilibrium by $q_n(t) = x_{n3}(t) - x_n^{(0)}$, and truncating at second order, the Lagrangian for the axial motion takes the form [219]

$$L = \frac{M}{2} \left[\sum_{m=1}^N \dot{q}_m^2 - \omega_z^2 \sum_{m,n=1}^N A_{mn} q_m q_n \right], \quad (1.6)$$

where the coupling matrix A_{mn} is given by

$$A_{mn} = \begin{cases} 1 + 2 \sum_{\substack{p=1 \\ p \neq m}}^N \frac{1}{|u_m - u_p|^3} & \text{if } m = n, \\ -\frac{2}{|u_m - u_n|^3} & \text{if } m \neq n. \end{cases} \quad (1.7)$$

Similarly, for the transverse directions ($i = 1, 2$), the displacements ξ_{ni} are governed by an analogous Lagrangian with coupling matrix [229]

$$B_{mn} = \left(\frac{1}{\alpha} + \frac{1}{2} \right) \delta_{mn} - \frac{1}{2} A_{mn}, \quad (1.8)$$

where $\alpha = (\omega_z/\omega_\perp)^2 \ll 1$ is the dimensionless anisotropy parameter characterizing the ratio of axial to transverse confinement, with $\omega_\perp$ denoting the transverse trapping frequency ($\omega_\perp = \omega_x$ or ω_y). The matrix B_{mn} shares the same eigenvectors as A_{mn} but has eigenvalues related to those of A_{mn} by [229]

$$\gamma_p = \frac{1}{\alpha} + \frac{1}{2} - \frac{\mu_p}{2}, \quad (1.9)$$

where μ_p are the eigenvalues of A_{mn} .

Since A_{mn} is real, symmetric, and positive definite, it possesses N non-negative eigenvalues μ_p ($p = 1, \dots, N$) with associated orthonormal eigenvectors $\mathbf{b}^{(p)}$, defined by

$$\sum_{n=1}^N A_{mn} b_n^{(p)} = \mu_p b_m^{(p)}, \quad p = 1, \dots, N, \quad (1.10)$$

with the orthonormality conditions $\sum_{n=1}^N b_n^{(p)} b_n^{(q)} = \delta_{pq}$. The eigenvectors are ordered by increasing eigenvalue. The first two modes are known analytically for any N [219]: the center-of-mass (COM) mode,

$$\mathbf{b}^{(1)} = \frac{1}{\sqrt{N}}(1, 1, \dots, 1), \quad \mu_1 = 1, \quad (1.11)$$

describing uniform translation of the entire crystal, and the stretch (or breathing) mode,

$$\mathbf{b}^{(2)} = \frac{1}{\sqrt{\sum_{m=1}^N u_m^2}}(u_1, u_2, \dots, u_N), \quad \mu_2 = 3, \quad (1.12)$$

in which each ion oscillates with an amplitude proportional to its equilibrium distance from the trap center. Higher modes must be determined numerically.

In terms of the normal-mode coordinates $Q_p(t) = \sum_{m=1}^N b_m^{(p)} q_m(t)$, the Lagrangian (1.6) decouples into N independent harmonic oscillators,

$$L = \frac{M}{2} \sum_{p=1}^N \left(\dot{Q}_p^2 - \nu_p^2 Q_p^2 \right), \quad (1.13)$$

with mode frequencies

$$\nu_p = \sqrt{\mu_p} \omega_z. \quad (1.14)$$

In particular, the COM mode has frequency $\nu_1 = \omega_z$ and the stretch mode has frequency $\nu_2 = \sqrt{3} \omega_z$. The transverse mode frequencies are given by $\Omega_p = \sqrt{\gamma_p} \omega_z$.

Phase transition to zigzag configuration

The linear chain configuration is stable only as long as all transverse mode frequencies Ω_p are real, which requires $\gamma_p > 0$ for all p . From Eq. (1.9), the transverse eigenvalues decrease with increasing μ_p , so the highest axial eigenvalue μ_N yields the lowest (and most critical) transverse eigenvalue. The condition $\gamma_N > 0$ gives [229]

$$\alpha < \alpha_{\text{crit}} = \frac{2}{\mu_N - 1}. \quad (1.15)$$

When the anisotropy parameter exceeds this critical value, the transverse frequency of the highest mode becomes imaginary, indicating a mechanical instability of the linear configuration. The ions then undergo a structural phase transition and arrange themselves in a two-dimensional zigzag pattern [230, 231]. This transition is a direct consequence of the competition between the transverse confinement, which favors a linear arrangement, and the Coulomb repulsion, which tends to push ions apart. As N increases, μ_N grows, and α_{crit} decreases, meaning that longer chains require tighter transverse confinement to remain stable.

Throughout this thesis we consider exclusively one-dimensional ion chains and always assume that the transversal trapping frequencies are chosen such that $\alpha \ll \alpha_{\text{crit}}$, ensuring the one-dimensional chain geometry.

Numerical calculation of phonon modes

The equilibrium positions and normal-mode frequencies enter the calculation of the effective spin-spin coupling strengths we derive in Sec. 1.2.1. In this subsection, we outline the numerical procedure used to determine these quantities.

For $N > 3$, the system of coupled nonlinear equations (1.4) must be solved numerically. Equivalently, one can minimize the dimensionless potential energy,

$$\tilde{V}(\{u_n\}) = \sum_{n=1}^N u_n^2 + \sum_{\substack{n,m=1 \\ m \neq n}}^N \frac{1}{|u_n - u_m|}, \quad (1.16)$$

which is obtained from Eq. (1.1) (restricted to the axial direction) by expressing the ion positions in units of ℓ . Both the gradient (1.4) and the fact that $\tilde{V}$ is bounded from below guarantee convergence of standard minimization algorithms.

In practice, we employ a gradient-based minimization (L-BFGS-B as implemented in `scipy.optimize.minimize`) with the analytically computed gradient, $\partial\hat{V}/\partial u_m$, which equals the left-hand side of Eq. (1.4) up to an overall factor of two. A suitable initial guess is constructed by distributing N ions uniformly over an interval whose extent is estimated from the empirical minimum spacing formula (1.5).

Once the equilibrium positions $\{u_n\}$ are known, the coupling matrix A_{mn} , Eq. (1.7), is constructed and diagonalized. This matrix is the (dimensionless) Hessian of the axial potential evaluated at equilibrium:

$$A_{mn} = \frac{1}{M\omega_z^2} \left. \frac{\partial^2 V}{\partial x_{m3} \partial x_{n3}} \right|_{\bar{x}}. \quad (1.17)$$

The eigenvalues μ_p and eigenvectors $b_n^{(p)}$ are obtained by a single symmetric matrix diagonalization (e.g., via the LAPACK routine `dsyev` as available through `scipy.linalg.eigh`). The axial mode frequencies then follow from Eq. (1.14) as $\nu_p = \sqrt{\mu_p} \omega_z$, and the transverse mode frequencies from $\Omega_p = \sqrt{\gamma_p} \omega_z$ with γ_p given by Eq. (1.9), without the need for a separate diagonalization of the transverse coupling matrix.

The resulting mode frequencies and eigenvectors constitute the numerical input for computing the effective spin-spin coupling matrix of the MAGIC scheme, as discussed in Sec. 1.2.2.

1.1.3 Typical noise sources

Trapped-ion qubits are generally considered one of the leading physical platforms for quantum information processing, with coherence times that can exceed seconds to minutes and gate fidelities at or beyond the fault-tolerance threshold [164, 168, 218, 228]. Nevertheless, a set of well-characterized technical noise sources ultimately limits the achievable fidelity and sets the requirements for the noise-mitigation strategy discussed in Chapter 4. Broadly, these sources fall into three categories [157, 217]: (i) qubit-frequency (dephasing) noise, driven by fluctuations of the external fields that define the qubit transition; (ii) control-field noise from the lasers or microwaves used to drive single- and two-qubit gates; and (iii) motional noise, arising from instabilities of the trap potential and from heating of the collective ion motion.

In the MAGIC scheme discussed in the next section, gates are driven by static magnetic-field gradients acting on ions encoded in magnetically sensitive hyperfine states. Field sensitivity is therefore unavoidable: the same gradient that enables the gate also makes the qubit splitting linearly dependent on the local magnetic field B through the Zeeman effect, so that fluctuations δB translate directly into fluctuations of the qubit frequency. This makes dephasing from ambient magnetic-field noise by far the dominant error channel any time magnetically sensitive qubit states are used, and motivates a more detailed discussion.

The magnetic field at the ion position in a typical laboratory environment contains a broad spectrum of noise components. Line-frequency noise at 50 Hz (or 60 Hz) and its harmonics, radiated by nearby electronic equipment, is usually the largest contribution, together with power-supply noise and instabilities of current sources and drivers in the tens-of-kilohertz range [232]. On long timescales, thermal drifts of coils or of permanent magnets add a slow-varying background that must be tracked through periodic recalibration. In an unshielded setup these effects combine to rms field fluctuations of $\gtrsim 100$ nT at fields of order 10 mT [232]. For a field-sensitive $^{171}\text{Yb}^+$ hyperfine qubit encoded in $|F=0, m_F=0\rangle \leftrightarrow |F=1, m_F=+1\rangle$, ~ 100 nT of rms field noise translates into qubit-frequency fluctuations of order 1 kHz, large enough to significantly degrade gate fidelity if left uncorrected. Standard

mitigation techniques include passive μ -metal shielding, active feedback on the coil currents using fluxgate sensors, which has been shown to suppress the rms noise to the few-nT level [232], and the replacement of electromagnets by temperature-stabilized permanent magnets [137].

1.2 Magnetic Gradient Induced Coupling (MAGIC)

The MAGIC scheme for quantum computation [161, 162, 233] uses microwaves instead of optical lasers to address ions and implement quantum gates. Since microwaves cannot be focused on individual ions, a magnetic field gradient is used to Zeeman-shift the resonance frequency of each ion and make them individually addressable in frequency space. Furthermore, the magnetic field gradient induces an Ising-type always-on all-to-all coupling between the ions, even in the absence of any driving fields, which can be used to synthesize quantum gates [234, 235] or perform adiabatic quantum sweeps [236]. In Sec. 1.2.1 we give a review of the derivation of the spin–spin coupling. In Sec. 1.2.2 we then provide concrete values of the achievable coupling strengths in the MAGIC setup for realistic experimental parameters.

1.2.1 Derivation of effective spin-spin coupling

We consider a one-dimensional (1D) ion chain, elongated in the z direction. The ions are subject to a static magnetic field with a constant gradient $\partial_z B$ in the axial direction z , $\mathbf{B}(z) = \mathbf{B}_0 + \partial_z B z \hat{\mathbf{e}}_z + \mathcal{O}(z^2)$, where $\mathbf{B}_0$ is a constant offset field. At the end of this section, we generalize this setting to gradients in the transversal directions. Moreover, we will discuss the effect of a field with non-vanishing curvature in Chapter 2. We further assume that we use magnetically sensitive states as our qubit states (typically states from a hyperfine ground state manifold with a resonance frequency in the microwave regime [163, 237, 238]) to make them individually addressable. Although such states are typically prone to dephasing caused by magnetic field fluctuations, long coherence times can be achieved via suitable protection schemes such as dynamical decoupling [119, 239] (see Chapter 4) or microwave dressing [240]. In Ref. [241], T_2 coherence times of 2.1(1) seconds were demonstrated for electron spin states in a single trapped $^{40}\text{Ca}^+$ ion using effective magnetic shielding techniques.

The spatially varying magnetic field gives rise to a position-dependent resonance frequency $\omega(z)$, which in general has to be calculated numerically or can be estimated using the Breit–Rabi formula [242]. Under weak Zeeman splitting, which holds for the examples considered in this work, provided that the offset field $\mathbf{B}_0$ is sufficiently small, the resonance frequency can be obtained as $\omega(z) = \omega_0 + (g_F^{(e)} m_F^{(e)} - g_F^{(g)} m_F^{(g)}) \mu_B B(z) / \hbar$. Here, $g_F^{(g)}$ and $g_F^{(e)}$ are, respectively, the Landé g -factors of the qubit ground and excited state with total atomic angular momentum F , $m_F^{(g)}$ and $m_F^{(e)}$ are the corresponding magnetic quantum numbers, and μ_B is the Bohr magneton. The bare resonance frequency in the absence of a magnetic field is denoted by ω_0 . The internal energy of the ions is then given by

$$H_{\text{int}} = -\frac{\hbar}{2} \sum_{n=1}^N \omega(z_n) \sigma_Z^{(n)}, \quad (1.18)$$

where $\sigma_Z^{(n)}$ denotes the Pauli- Z operator acting on qubit n at position z_n and N is the number of ions in the chain.

The resonance frequency can be expanded around the equilibrium position $z_{n,0}$ of ion n

as

$$\omega(z_n) = \omega_n + \partial_z \omega_n (z_n - z_{n,0}) + \mathcal{O}(z_n^2), \quad (1.19)$$

where $\omega_n = \omega(z_{n,0})$ is the resonance frequency and $\partial_z \omega_n = \partial \omega(z_{n,0}) / \partial z$ the first derivative at the position of the ion. Since the excursions $q_n = z_n - z_{n,0}$ from the equilibrium positions are small compared to the inter-ion spacing (typically on the order of the oscillator length $\Delta z \sim \text{nm}$, while the inter-ion distance is $\sim \mu\text{m}$), we can truncate this expansion at linear order. Higher-order terms in this expansion will be discussed in Chapter 2. We can then write the internal energy as

$$H_{\text{int}} = -\frac{\hbar}{2} \sum_{n=1}^N \omega_n \sigma_Z^{(n)} - \frac{\hbar}{2} \sum_{n=1}^N \partial_z \omega_n q_n \sigma_Z^{(n)}. \quad (1.20)$$

Additionally, the ions have an energy associated with their external degrees of freedom, composed of their kinetic energy, the potential energy due to the trap, as well as their mutual Coulomb repulsion:

$$\begin{aligned} H_{\text{ext}} &= H_{\text{kin}} + V(\mathbf{z}) \\ &= H_{\text{kin}} + V_{\text{trap}}(\mathbf{z}) + V_{\text{Coul}}(\mathbf{z}) \\ &= \frac{1}{2m} \sum_{n=1}^N p_{z,n}^2 + \frac{m}{2} \omega_z^2 \sum_{n=1}^N z_n^2 + \frac{e^2}{8\pi\epsilon_0} \sum_{n \neq l}^N \frac{1}{|z_n - z_l|}. \end{aligned} \quad (1.21)$$

Here, the vector $\mathbf{z} = (z_1, \dots, z_N)^T$ contains the axial positions of all ions, e is the charge of the ions, m their mass, ϵ_0 the vacuum permittivity, and $p_{z,n}$ is the momentum of ion n in the axial direction. For now, we will consider a harmonic trap with angular frequency ω_z . The equilibrium positions $\mathbf{z}_0$ are determined by the condition $\partial V(\mathbf{z}) / \partial z_n|_{\mathbf{z}_0} = 0$, i.e., the balance of the trapping force and the mutual Coulomb repulsion. Expanding the potential energy around this minimum, the linear terms vanish by definition and, after dropping a constant energy offset, we obtain

$$H_{\text{ext}} = \frac{1}{2m} \sum_{n=1}^N p_{z,n}^2 + \frac{m}{2} \sum_{n=1}^N \sum_{l=1}^N A_{nl} q_n q_l + \mathcal{O}(q^3), \quad (1.22)$$

with $A_{nl} = \partial^2 V(\mathbf{z}_0) / \partial z_n \partial z_l$. In Chapter 2, we will discuss the effects of higher-order terms in the expansion of the potential energy.

The matrix A is symmetric and real and can thus be diagonalized as $D = S^{-1} A S$ with an orthogonal matrix S . Since A is a positive-definite Hessian (it results from an expansion around a minimum of the potential energy), its eigenvalues are positive and can be written as ν_l^2 . We can now express the local coordinates q_n in terms of collective normal modes Q_l ,

$$q_n = \sum_{l=1}^N S_{nl} Q_l. \quad (1.23)$$

The normal modes Q_l describe collective oscillation patterns of the ion chain, with the matrix elements S_{nl} quantifying the participation of ion n in mode l . The lowest-frequency mode ($l = 1$, the center-of-mass mode) corresponds to all ions oscillating in phase with $S_{n,1} = 1/\sqrt{N}$, while higher modes involve increasingly complex oscillation patterns with nodes along the chain.

These normal modes diagonalize the external Hamiltonian to

$$H_{\text{ext}} = \frac{1}{2m} \sum_{n=1}^N P_n^2 + \frac{m}{2} \sum_{n=1}^N \nu_n^2 Q_n^2, \quad (1.24)$$

where we have defined the normal-mode momenta $P_n = \sum_{l=1}^N (S^{-1})_{nl} p_{z,l}$ and only written down the second-order terms.

We can quantize the external Hamiltonian by introducing phonon creation and annihilation operators for the normal modes:

$$Q_l = \Delta z_l (a_l^\dagger + a_l), \quad (1.25)$$

with $\Delta z_l = \sqrt{\hbar/2m\nu_l}$ being the width of the ground-state wave packet of the quantum harmonic oscillator. In terms of these operators, the external Hamiltonian becomes $H_{\text{ext}} = \hbar \sum_{n=1}^N \nu_n a_n^\dagger a_n$ and the overall Hamiltonian $H = H_{\text{int}} + H_{\text{ext}}$ can now be expressed as

$$\begin{aligned} H = & -\frac{\hbar}{2} \sum_{n=1}^N \omega_n \sigma_Z^{(n)} - \frac{\hbar}{2} \sum_{n=1}^N \sum_{l=1}^N \partial_z \omega_n S_{nl} \Delta z_l (a_l^\dagger + a_l) \sigma_Z^{(n)} \\ & + \hbar \sum_{n=1}^N \nu_n a_n^\dagger a_n. \end{aligned} \quad (1.26)$$

We can decouple the internal and external degrees of freedom (spins and phonons) via a suitable unitary transformation [161, 233] (in some contexts known as the polaron transformation [243]), $\tilde{H} = U^\dagger H U$, with

$$U = \exp \left(-\frac{i}{\hbar} \sum_{n=1}^N \sum_{l=1}^N \Delta z_l \epsilon_{nl} \sigma_Z^{(n)} P_l \right), \quad (1.27a)$$

$$\epsilon_{nl} = \frac{\Delta z_l \partial_z \omega_n}{\nu_l} S_{nl}. \quad (1.27b)$$

Here, we have defined the effective Lamb–Dicke parameter ϵ_{nl} , describing how strongly ion n (with internal transition frequency ω_n) couples to phonon mode l (with frequency ν_l).

With the transformed phonon operators

$$U^\dagger a_l U = a_l + \frac{1}{2} \sum_n \epsilon_{nl} \sigma_Z^{(n)} \quad (1.28)$$

(and unchanged spin operators $\tilde{\sigma}_Z^{(n)} = \sigma_Z^{(n)}$), we can now derive the transformed Hamiltonian explicitly. Substituting Eq. (1.28) into the phonon part of H , we obtain

$$\hbar \sum_l \nu_l \left(a_l + \frac{1}{2} \sum_n \epsilon_{nl} \sigma_Z^{(n)} \right)^\dagger \left(a_l + \frac{1}{2} \sum_m \epsilon_{ml} \sigma_Z^{(m)} \right). \quad (1.29)$$

Expanding this product yields three types of terms: the original phonon Hamiltonian $\hbar \sum_l \nu_l a_l^\dagger a_l$, terms linear in $a_l^\dagger$ and a_l , and a term quadratic in the spin operators. The terms linear in the phonon operators are

$$\frac{\hbar}{2} \sum_{n=1}^N \sum_{l=1}^N \nu_l \epsilon_{nl} (a_l^\dagger + a_l) \sigma_Z^{(n)}, \quad (1.30)$$

which, using the definition of ϵ_{nl} in Eq. (1.27b), exactly cancel the spin–phonon coupling in H_{int} [second term of Eq. (1.20)]. This cancellation is the purpose of the polaron transformation and holds by construction for the specific choice of ϵ_{nl} . The total term quadratic in the spin operators (adding all terms from the total Hamiltonian) reads

$$-\frac{\hbar}{4} \sum_l \nu_l \sum_{n,m} \epsilon_{nl} \epsilon_{ml} \sigma_Z^{(n)} \sigma_Z^{(m)}. \quad (1.31)$$

Separating this into the diagonal ($n = m$) and off-diagonal ($n \neq m$) contributions, the diagonal terms yield a constant energy shift since $(\sigma_Z^{(n)})^2 = 1$, while the off-diagonal terms produce the desired pairwise spin–spin interaction. The total Hamiltonian thus takes the form

$$\begin{aligned} \tilde{H}^{(2)} = & -\frac{\hbar}{2} \sum_{n=1}^N \omega_n \sigma_Z^{(n)} + \hbar \sum_{n=1}^N \nu_n a_n^\dagger a_n \\ & -\frac{\hbar}{2} \sum_{i < j}^N J_{ij}^{(2)} \sigma_Z^{(i)} \sigma_Z^{(j)}, \end{aligned} \quad (1.32)$$

with the spin–spin coupling strength

$$J_{ij}^{(2)} = \sum_{n=1}^N \nu_n \epsilon_{in} \epsilon_{jn}. \quad (1.33)$$

The effective Hamiltonian $\tilde{H}^{(2)}$ relies on two approximations made *prior* to the polaron transformation: (i) the expansion of the potential energy to second order in the displacements q_n , and (ii) the expansion of the resonance frequency to first order in q_n . The polaron transformation itself is exact, i.e., it introduces no further approximations beyond these truncations. Higher-order terms in both expansions will be the subject of the following chapter.

So far, we have restricted the derivation to the axial degree of freedom. In a real three-dimensional trap, however, the ions also possess two transversal motional degrees of freedom (perpendicular to the chain axis), each with their own set of collective normal modes (see discussion in Sec. 1.1.2). At second order in the expansion of the potential energy, the axial and transversal modes decouple due to the cylindrical symmetry of the linear equilibrium configuration [229]. Each spatial direction $\alpha = 1, 2, 3$ can therefore be independently diagonalized into normal modes with frequencies $\nu_n^{(\alpha)}$ and mode participation matrices $S^{(\alpha)}$. If the applied magnetic field has a non-vanishing gradient also in the transversal directions, the polaron transformation can be carried out independently for each set of modes, yielding a generalized spin–spin coupling strength

$$J_{ij}^{(2)} = \sum_{\alpha=1}^3 \sum_{n=1}^N \nu_n^{(\alpha)} \epsilon_{in}^{(\alpha)} \epsilon_{jn}^{(\alpha)}, \quad (1.34)$$

where $\epsilon_{in}^{(\alpha)}$ is the effective Lamb–Dicke parameter describing the coupling of ion i to the n -th mode in direction α , defined analogously to Eq. (1.27b). Each direction contributes proportionally to $(\partial_\alpha B / \nu^{(\alpha)})^2$, where $\partial_\alpha B$ is the magnetic field gradient in direction α . In the standard MAGIC setup, where the gradient is applied only along the axial direction and vanishes in both transversal directions, the transversal Lamb–Dicke parameters $\epsilon_{in}^{(\alpha)}$ are zero

for $\alpha = 1, 2$, and Eq. (1.34) reduces to the purely axial expression in Eq. (1.33). The effects of transversal modes at third order in the potential energy, which can couple to the axial modes through Coulomb anharmonicities, will be discussed in Chapter 2.

The spin–spin coupling strength $J_{ij}^{(2)}$ can be tuned via the magnetic field gradient and the strength of the external harmonic trap, since $J_{ij}^{(2)} \sim (\partial_z B / \nu)^2$. By employing different trapping potentials along the chain (not necessarily harmonic), coupling strengths between different pairs of ions can be enhanced or reduced [223]. This tunability is a key feature of the MAGIC platform, as the effective Hamiltonian $\tilde{H}^{(2)}$ in Eq. (1.32) directly realizes a long-range Ising model $H_{\text{Ising}} = -\sum_{i < j} J_{ij} \sigma_Z^{(i)} \sigma_Z^{(j)} - \sum_i h_i \sigma_Z^{(i)}$, and thus becomes interesting for quantum simulation experiments [244–246]. At the same time, this Hamiltonian structure is central to quantum optimization, where a combinatorial optimization problem is encoded in a cost function of the form $C(\mathbf{s}) = \sum_{i < j} J_{ij} s_i s_j + \sum_i h_i s_i$ with classical spin variables $s_i = \pm 1$ [64, 189]. By mapping this cost function onto the coupling matrix $J_{ij}^{(2)}$ and the local fields, the ground state of the Ising Hamiltonian corresponds to the optimal solution. The ability to engineer the coupling strengths, for instance via segmented trapping potentials, is therefore essential for encoding arbitrary problem instances into the native interactions of the ion chain. The application of this native spin–spin coupling for quantum annealing will be the subject of Chapters 3 and 4.

1.2.2 Numerical example with $^{171}\text{Yb}^+$ ions

To illustrate the magnitude of the spin–spin coupling strengths achievable in practice, we now evaluate Eq. (1.33) numerically for concrete experimental parameters. The computation proceeds by first determining the equilibrium positions of the ions (as described in Sec. 1.1.2), then diagonalizing the Hessian of the potential energy to obtain the normal mode frequencies ν_n and mode participation matrix S_{nl} , and finally evaluating the coupling strength $J_{ij}^{(2)} = \sum_n \nu_n \epsilon_{in} \epsilon_{jn}$ from the resulting Lamb–Dicke parameters.

As a first example, we consider a chain of five $^{171}\text{Yb}^+$ ions in a harmonic trap with axial frequency $\omega_z = 2\pi \times 130 \text{ kHz}$ and a magnetic field gradient of $\partial_z B = 19 \text{ T m}^{-1}$, as realized experimentally in Ref. [163]. We choose as qubit states $|F=0, m_F=0\rangle$ and $|F=1, m_F=+1\rangle$ ($g_F \approx 1$) from the ground state hyperfine manifold. For this choice, the difference of Landé factors and magnetic quantum numbers entering the general expression for the resonance frequency ($\omega(z) = \omega_0 + (g_F^{(e)} m_F^{(e)} - g_F^{(g)} m_F^{(g)}) \mu_B B(z) / \hbar$, using the Breit–Rabi formula [242]) reduces to $g_F^{(e)} m_F^{(e)} - g_F^{(g)} m_F^{(g)} = 1 \cdot (+1) - 0 = 1$, so that the resonance frequency simplifies to $\omega(z) \simeq \omega_0 + \mu_B B(z) / \hbar$ [235].

It is important to note that the equilibrium positions of ions in a harmonic trap are not equally spaced [219]. As a consequence, the coupling strength $J_{ij}^{(2)}$ between neighboring ions is not uniform along the chain, and the strongest coupling is found between the two outermost neighboring pairs. This results from the spatial structure of the higher-order phonon modes entering Eq. (1.33), whose mode participation S_{nl} is largest near the edges of the chain. For the above parameters, the largest coupling strength is $J_{12}^{(2)} = J_{45}^{(2)} \approx 2\pi \times 26.5 \text{ Hz}$, occurring between the two outermost neighboring pairs. The weakest coupling, $J_{15}^{(2)} \approx 2\pi \times 12.9 \text{ Hz}$, occurs between the two most distant ions at opposite ends of the chain. For a stronger magnetic field gradient of 150 T m^{-1} , which is feasible with current technology [238, 247], the maximal (minimal) spin–spin coupling strength increases to $J_{12}^{(2)} = J_{45}^{(2)} \approx 2\pi \times 1650 \text{ Hz}$ ($J_{15}^{(2)} \approx 2\pi \times 805 \text{ Hz}$).

In Fig. 1.1, we plot the scaling of the maximal and minimal coupling strengths with




Figure 1.1: Spin–spin coupling strengths in the MAGIC scheme, calculated as a function of the number of $^{171}\text{Yb}^+$ ions N in a chain with trapping frequency $\omega_z = 2\pi \times 130$ kHz and magnetic field gradient (a) 150 T m^{-1} and (b) 19 T m^{-1} . Depicted are only the largest coupling strength $J_{\text{max}}^{(2)}$ (blue), occurring between two neighboring ions at the edges of the chain, as well as the smallest coupling strength $J_{\text{min}}^{(2)}$ (orange), felt by the most distant ions at opposite ends of the chain. The decrease of $J_{\text{max}}^{(2)}$ follows the law $N^{-1/2}$ to good approximation, while $J_{\text{min}}^{(2)}$ decreases approximately as $N^{-1.19}$. Corresponding fits are shown as grey dashed lines.

the number of ions N in the chain, for the same trapping frequency and for both choices of magnetic field gradient. From a numerical fit (grey dashed lines), the maximal spin–spin coupling strength decreases as $N^{-0.51 \pm 0.01}$ ¹, matching the $N^{-1/2}$ scaling known from the Cirac–Zoller gate [158] and the Mølmer–Sørensen gate [248], which share the same fundamental origin in the phonon-mediated nature of the interaction. The minimal spin–spin coupling, felt by the two ions at opposite ends of the chain, diminishes more rapidly as $N^{-1.19 \pm 0.01}$. This stronger decrease can be attributed to the large distance between the involved ions (scaling linearly with N) in interplay with contributions from phonon modes beyond the center-of-mass mode, which can give rise to spatially decreasing interactions [201, 245, 249–251].

The dependence of the coupling strength on the ion species enters through the oscillator length $\Delta z_l = \sqrt{\hbar/2m\nu_l}$ appearing in the Lamb–Dicke parameter [Eq. (1.27b)]. Since

¹Uncertainties in the fit parameters are computed from the estimated covariance matrix returned by the `curve_fit` function from the SciPy library. The underlying data points are calculated up to numerical precision, i.e., without considering any errors in the physical parameters.

$\Delta z_l \propto 1/\sqrt{m}$ and $J_{ij}^{(2)}$ is quadratic in $\epsilon_{nl} \propto \Delta z_l$, the two-body coupling scales as $J_{ij}^{(2)} \propto 1/m$. Lighter ions therefore achieve stronger spin–spin couplings and correspondingly faster quantum gate speeds. In Tab. 2.1 in the next chapter, we compare numerical values of the two-body coupling strengths for $^{171}\text{Yb}^+$ and $^9\text{Be}^+$ ion chains of different lengths and for both magnetic field gradients; this table additionally contains estimates of higher-order corrections, which will be discussed in the following chapter. For the $^9\text{Be}^+$ ions, we consider the qubit states $|F=1, m_F=0\rangle$ and $|F=2, m_F=+1\rangle$ from the hyperfine ground state manifold. The numerical values confirm the $1/m$ scaling, with $^9\text{Be}^+$ exhibiting coupling strengths roughly a factor of $m_{\text{Yb}}/m_{\text{Be}} = 19$ larger than $^{171}\text{Yb}^+$. However, certain isotopes of heavier ion species can have other advantages that counterbalance the slower gate times. In particular, $^{171}\text{Yb}^+$ ions have a simple hyperfine level structure due to their nuclear spin of $I = 1/2$ (in contrast to $I = 3/2$ for $^9\text{Be}^+$), which can be advantageous for state preparation and cooling. The MAGIC platform developed by eleQtron and considered throughout this work employs $^{171}\text{Yb}^+$ ions.

1.3 Implementing discrete quantum gates with trapped ions

A universal gate set for quantum computation requires the ability to perform arbitrary single-qubit rotations together with at least one entangling two-qubit gate [252]. In trapped-ion systems, both types of operations exploit the coupling between the ions’ internal (pseudo-spin) states and their shared motional degrees of freedom (phonons), mediated by external driving fields or, in the case of the MAGIC platform, by a static magnetic field gradient. In the following, we first review how single-qubit rotations are realized through resonant driving on the carrier transition (Sec. 1.3.1). We then discuss the two paradigmatic proposals for entangling gates on trapped-ions platforms: the Cirac–Zoller gate [158], which uses a single shared phonon as a quantum data bus (Sec. 1.3.2), and the Mølmer–Sørensen gate [159, 160, 248], which generates entanglement through geometric phases in motional phase space (Sec. 1.3.2). Finally, we describe how entangling operations are implemented on the MAGIC platform, where the magnetic-field-gradient-induced spin–spin coupling enables native ZZ gates without any driving fields, and where the equivalence of the microwave driving Hamiltonian to the standard light–matter interaction permits the use of the same driven gate schemes as in laser-based setups (Sec. 1.3.3).

1.3.1 Single-qubit gates

We consider a string of N ions confined in a linear Paul trap, with collective motional (normal) modes at frequencies ω_m ($m = 1, \dots, N$). Each ion encodes a two-level qubit $\{|\downarrow\rangle, |\uparrow\rangle\}$ in two of its internal states with transition frequency ω_0 . A driving field (laser or microwave) of frequency ω_d , phase ϕ , and Rabi frequency Ω couples to ion i . In the interaction picture with respect to the free Hamiltonian $H_0 = \frac{\hbar\omega_0}{2}\sigma_z^{(i)} + \sum_m \hbar\omega_m \hat{a}_m^\dagger \hat{a}_m$, and after the rotating-wave approximation, the interaction Hamiltonian takes the form

$$H_{\text{int}} = \frac{\hbar\Omega}{2} \sigma_+^{(i)} \exp[i\eta_{im}(\hat{a}_m e^{-i\omega_m t} + \hat{a}_m^\dagger e^{i\omega_m t})] e^{-i(\Delta t + \phi)} + \text{H.c.}, \quad (1.35)$$

where $\Delta = \omega_d - \omega_0$ is the detuning from the qubit transition and $\eta_{im} = S_{im} k \sqrt{\hbar/2m\omega_m}$ is the Lamb–Dicke parameter of ion i on mode m , with S_{im} the normal-mode eigenvector element, k the effective wavevector projected onto the mode axis, and m the ion mass. For

notational clarity we display the coupling to a single motional mode m . The multi-mode form is recovered by replacing $\eta_{im}(\hat{a}_m e^{-i\omega_m t} + \hat{a}_m^\dagger e^{i\omega_m t})$ in the exponent with $\sum_m \eta_{im}(\hat{a}_m e^{-i\omega_m t} + \hat{a}_m^\dagger e^{i\omega_m t})$; since the Lamb-Dicke expansion below is linear in the mode operators, each mode contributes independently and the carrier and sideband analysis is unaffected. Physically, η_{im} is the ratio of the motional ground-state wavepacket size $x_0 = \sqrt{\hbar/2M\omega_m}$ to the reduced wavelength $1/k$ of the driving field. The system is in the Lamb-Dicke regime when $\eta^2(2\bar{n}+1) \ll 1$. In this regime, the exponential in Eq. (1.35) can be expanded to first order,

$$\exp[i\eta_{im}(\hat{a}_m e^{-i\omega_m t} + \hat{a}_m^\dagger e^{i\omega_m t})] \approx 1 + i\eta_{im}(\hat{a}_m e^{-i\omega_m t} + \hat{a}_m^\dagger e^{i\omega_m t}), \quad (1.36)$$

which separates the interaction into three spectral components, each selected by the choice of detuning Δ .

For single-qubit operations, the driving field is tuned to the carrier transition ($\Delta \approx 0$), where the interaction drives spin flips without changing the motional state. In this case the phonon operators in Eq. (1.36) contribute only off-resonant terms (detuned by $\pm\omega_m \gg \Omega$) that can be neglected, and the Hamiltonian reduces to

$$H_{\text{car}} = \frac{\hbar\Omega}{2}(e^{i\phi}\sigma_+^{(i)} + e^{-i\phi}\sigma_-^{(i)}) + \frac{\hbar\Delta}{2}\sigma_z^{(i)}. \quad (1.37)$$

On exact resonance ($\Delta = 0$), this simplifies to

$$H_{\text{car}} = \frac{\hbar\Omega}{2}(\cos\phi\sigma_x^{(i)} - \sin\phi\sigma_y^{(i)}), \quad (1.38)$$

which generates a rotation around the axis $\hat{n} = (\cos\phi, -\sin\phi, 0)$ on the Bloch sphere with rate Ω . By choosing the phase ϕ and the pulse duration τ , arbitrary rotations $R_{\hat{n}}(\theta) = \exp(-i\theta\hat{n} \cdot \boldsymbol{\sigma}/2)$ with $\theta = \Omega\tau$ around any axis in the XY plane can be implemented. Rotations around the Z axis are typically realized virtually by updating the phase reference of subsequent pulses [253], avoiding the need for a physical operation and its associated errors.

The remaining two spectral components of the Lamb-Dicke expansion (1.36) correspond to the red sideband ($\Delta = -\omega_m$), coupling $|\downarrow, n\rangle \leftrightarrow |\uparrow, n-1\rangle$ at Rabi frequency $\eta\Omega\sqrt{n}$, and the blue sideband ($\Delta = +\omega_m$), coupling $|\downarrow, n\rangle \leftrightarrow |\uparrow, n+1\rangle$ at Rabi frequency $\eta\Omega\sqrt{n+1}$. These motional sideband transitions, which simultaneously flip the spin and change the phonon number by one, form the basis of the two-qubit entangling gate schemes discussed in the following sections.

1.3.2 Two-qubit entangling gates

A universal gate set for quantum computation requires, in addition to arbitrary single-qubit rotations, at least one entangling two-qubit gate [252]. In trapped-ion systems, all high-fidelity entangling gates exploit the shared collective motional (phonon) degrees of freedom of ions in a common trap potential to mediate effective spin-spin interactions. We review the two paradigmatic proposals: the Cirac-Zoller gate [158], which uses a single phonon as a quantum data bus, and the Mølmer-Sørensen gate [159, 160, 248], which generates entanglement through spin-dependent forces and geometric phases in motional phase space.

The Cirac-Zoller gate

The proposal of Cirac and Zoller [158] was the first concrete scheme for universal quantum computation with trapped ions. It operates on a pair of ions A and B sharing a collective

motional mode (typically the center-of-mass mode at frequency ω_{COM}), which must be prepared in its ground state $|n = 0\rangle$.

The Cirac–Zoller CNOT gate is realized by a sequence of three individually addressed laser pulses [158, 254]:

1. A π -pulse on the red sideband of ion A maps the qubit state onto the phonon bus mode: $|\uparrow_A, 0\rangle \rightarrow -i|\downarrow_A, 1\rangle$, while $|\downarrow_A, 0\rangle$ is unaffected since the red sideband coupling vanishes at $n = 0$.
2. A 2π -pulse on the red sideband of ion B drives the transition $|\downarrow\rangle \leftrightarrow |a\rangle$ to an auxiliary level $|a\rangle$. The state $|\downarrow_B, 1\rangle$ undergoes a full Rabi oscillation through $|a_B, 0\rangle$ and back, acquiring a sign change: $|\downarrow_B, 1\rangle \rightarrow -|\downarrow_B, 1\rangle$. States with zero phonons are unaffected.
3. A second π -pulse on the red sideband of ion A reverses the mapping of Step 1, restoring the bus to the vacuum.

The net operation on the two-qubit computational basis is a controlled-Z gate, $|\uparrow\uparrow\rangle \rightarrow -|\uparrow\uparrow\rangle$ with all other basis states unchanged, which can be converted to a CNOT by single-qubit Hadamard gates on the target ion. Schmidt-Kaler *et al.* [254] demonstrated this gate on two $^{40}\text{Ca}^+$ ions with a fidelity of approximately 71%.

The principal limitation of the Cirac–Zoller scheme is its strict requirement for ground-state cooling of the phonon bus mode ($\bar{n} \approx 0$). Due to the n -dependent sideband Rabi frequencies $\eta\Omega\sqrt{n}$, the gate is highly sensitive to residual thermal excitation and motional heating during the pulse sequence.

The Mølmer-Sørensen gate

The Mølmer-Sørensen (MS) gate [159, 160, 248, 255] avoids the above limitation for strict ground-state cooling by generating entanglement through a spin-dependent force that drives state-dependent displacements in motional phase space, with the motional degrees of freedom returning to their initial state at the end of the gate.

The scheme applies a bichromatic laser field (two different frequencies) symmetrically detuned from the qubit resonance frequency by $\pm(\omega_m + \delta)$, where ω_m is the frequency of a collective motional mode and $\delta \ll \omega_m$ is a small detuning from the sidebands. The combined effect of the near-resonant red and blue sideband excitations, in the interaction picture and Lamb–Dicke regime, gives rise to the following Hamiltonian [256]:

$$H(t) = \frac{1}{2} \sum_{i,m} \eta_{im} \Omega_i \hat{\sigma}_\phi^{(i)} [\hat{a}_m^\dagger e^{-i\delta_m t} + \hat{a}_m e^{i\delta_m t}], \quad (1.39)$$

where $\delta_m = \mu - \omega_m$ is the detuning of the bichromatic beatnote μ from mode m , η_{im} is the Lamb–Dicke parameter of ion i on mode m , Ω_i is the Rabi frequency, and $\hat{\sigma}_\phi^{(i)} = \sigma_x^{(i)} \cos \phi + \sigma_y^{(i)} \sin \phi$ is a spin operator acting on ion i , with the angle ϕ determined by the optical phases of the driving fields.

The structure of Eq. (1.39) is that of a driven harmonic oscillator, with the driving amplitude conditioned on the spin state. To see this, recall that the displacement operator for a single bosonic mode is defined as

$$\hat{D}(\alpha) = \exp(\alpha \hat{a}^\dagger - \alpha^* \hat{a}), \quad (1.40)$$

which translates any motional state in phase space: $\hat{D}(\alpha) |\beta\rangle = e^{i\text{Im}(\alpha\beta^*)} |\alpha + \beta\rangle$ for a coherent state $|\beta\rangle$, and in particular $\hat{D}(\alpha) |0\rangle = |\alpha\rangle$. Crucially, the displacement is independent of

the initial state: a thermal mixture is translated by the same amplitude α as the vacuum. A Hamiltonian that is linear in $\hat{a}$ and $\hat{a}^\dagger$, of the form $H = f(t)\hat{a}^\dagger + f^*(t)\hat{a}$, generates precisely such displacements under time evolution. In Eq. (1.39), the prefactor $\hat{\sigma}_\phi^{(i)}$ makes the driving amplitude spin-dependent: eigenstates of $\hat{\sigma}_\phi^{(i)}$ with eigenvalues ± 1 experience opposite forces, and hence opposite displacements in the phase space of each motional mode. This spin-dependent displacement is the mechanism by which the MS gate converts motional excitations into spin-spin entanglement.

Since $H(t)$ is linear in the bosonic operators within the Lamb-Dicke approximation, it can be shown that the Magnus expansion terminates exactly at second order [256, 257], yielding the propagator

$$\hat{U}(\tau) = \exp \left[\sum_i \hat{\zeta}_i(\tau) \hat{\sigma}_\phi^{(i)} - i \sum_{i < j} \chi_{ij}(\tau) \hat{\sigma}_\phi^{(i)} \hat{\sigma}_\phi^{(j)} \right], \quad (1.41)$$

with two physically distinct contributions. The first-order term $\hat{\zeta}_i = \sum_m [\alpha_{im} \hat{a}_m^\dagger - \alpha_{im}^* \hat{a}_m]$ generates spin-dependent displacements in motional phase space, where the displacement amplitude for ion i on mode m is

$$\alpha_{im}(\tau) = \frac{\eta_{im} \Omega_i}{2\delta_m} (1 - e^{-i\delta_m \tau}). \quad (1.42)$$

This traces a circular trajectory of radius $|\alpha_{\max}| = \eta_{im} \Omega_i / (2|\delta_m|)$ in the phase space of mode m . The second-order term yields the spin-spin entangling phase

$$\chi_{ij}(\tau) = \sum_m \frac{\eta_{im} \eta_{jm} \Omega_i \Omega_j}{\delta_m} \left(\tau - \frac{\sin \delta_m \tau}{\delta_m} \right). \quad (1.43)$$

For the gate to produce a pure spin-spin entangling operation, the motional modes must disentangle from the spins at the gate time τ_g . From Eq. (1.42), the condition $\alpha_{im}(\tau_g) = 0$ for all i, m requires

$$\delta_m \tau_g = 2\pi k_m, \quad k_m \in \mathbb{Z}^+. \quad (1.44)$$

When this condition is satisfied, the operator $\hat{\zeta}_i$ vanishes and the propagator reduces to

$$\hat{U}(\tau_g) = \exp \left[-i \sum_{i < j} \chi_{ij}(\tau_g) \hat{\sigma}_\phi^{(i)} \hat{\sigma}_\phi^{(j)} \right]. \quad (1.45)$$

A maximally entangling gate, equivalent in entangling power to a CNOT, is obtained when $\chi_{ij} = \pi/4$. For $\phi = 0$, this gives $\hat{U} = \exp(-i\frac{\pi}{4}\sigma_x\sigma_x)$, which maps $|\downarrow\downarrow\rangle$ to the Bell state $(|\downarrow\downarrow\rangle - i|\uparrow\uparrow\rangle)/\sqrt{2}$.

The entangling phase χ_{ij} has a geometric interpretation [255]: it equals the area enclosed by the phase-space trajectory of the spin-dependent displacement. Since the displacement operator $\hat{D}(\alpha)$ shifts any initial state by the same amplitude α regardless of the phonon number, the enclosed area, and hence χ_{ij} , is independent of the initial motional occupation. This renders the MS gate insensitive to the initial thermal state of the motional modes, requiring only that the Lamb-Dicke approximation $\eta^2(2\bar{n} + 1) \ll 1$ holds. The MS gate furthermore does not require individual addressing: a single pair of globally applied bichromatic beams drives the entangling interaction on all ions simultaneously. These advantages over the Cirac-Zoller scheme have made geometric phase gates, whether in the

Mølmer-Sørensen ($\hat{\sigma}_\phi \hat{\sigma}_\phi$) or light-shift ($\hat{\sigma}_z \hat{\sigma}_z$) variant [216, 255], the standard approach for high-fidelity entangling operations in laser-based trapped-ion quantum processors, with state-of-the-art fidelities exceeding 99.9% [164, 165] and recently 99.99% using electronically controlled geometric phase gates [168].

1.3.3 Implementing gates with MAGIC

On the MAGIC platform, the magnetic-field-gradient-induced spin-spin coupling derived in the previous section [cf. Eq. (1.32)] is an always-on ZZ interaction. This interaction can directly be used to implement entangling two-qubit gates without any driving fields. To realize a unitary ZZ rotation of the form $\exp(-i\theta \sigma_Z^{(i)} \sigma_Z^{(j)})$ between ions i and j , one first transfers all other ions into internal states that are, to first order, magnetically insensitive, effectively decoupling them from the interaction. The Hamiltonian then reduces to

$$\tilde{H}^{(2)} = -\frac{\hbar}{2} J_{ij}^{(2)} \sigma_Z^{(i)} \sigma_Z^{(j)}, \quad (1.46)$$

and free evolution for a time T implements the desired gate with rotation angle $\theta = -\frac{1}{2} J_{ij}^{(2)} T$. More complex multi-qubit entangling gates can be synthesized by selectively hiding and unhiding subsets of ions and letting the system evolve freely in an alternating fashion [235].

Since the microwave driving Hamiltonian on the MAGIC platform, Eq. (1.35), takes the same form as the standard light-matter interaction Hamiltonian in laser-based trapped-ion setups [162] (substituting in the effective Lamb-Dicke parameter defined in Eq. (1.27b)), the full repertoire of driven quantum gates can equally be implemented using microwave radiation in place of laser fields. In particular, this includes single-qubit rotations via the carrier transition [Eq. (1.37)], as discussed in Sec. 1.3.1, as well as driven entangling gates such as the Mølmer-Sørensen gate [159, 160, 248]. An advantage of microwave control is the elimination of spontaneous photon scattering, which constitutes a fundamental error source in laser-driven gates [258], together with the high intrinsic stability of microwave sources in frequency, phase, and amplitude [259, 260].

A relevant source of systematic error in always-on ZZ gates on the MAGIC platform are undesired longitudinal fields of the form $h_Z \sigma_Z$, which shift the ions' resonance frequencies and introduce additional unitary contributions $\exp(-i h_Z T \sigma_Z)$ to the gate operation. As discussed in detail in Ref. [205] and the next chapter, such terms arise naturally from higher-order contributions in the MAGIC scheme, in particular from phonon-occupation-dependent corrections to the local fields whose strength increases with chain length and can become comparable to the two-body coupling $J_{ij}^{(2)}$ itself.

To illustrate the impact of such errors, consider the preparation of the Bell state $(|00\rangle + |11\rangle)/\sqrt{2}$ via the circuit $\text{CNOT}_{1,2} H_1$ applied to $|00\rangle$, where the CNOT gate is decomposed into a ZZ rotation and single-qubit operations as [234]

$$\text{CNOT}_{1,2} = e^{i(\pi/4)\sigma_Z^{(1)}} e^{i(\pi/4)\sigma_X^{(2)}} H_2 e^{-i(\pi/4)\sigma_Z^{(1)}\sigma_Z^{(2)}} H_2. \quad (1.47)$$

If the ZZ rotation is contaminated by a spurious local field,

$$e^{-i(\pi/4)\sigma_Z^{(1)}\sigma_Z^{(2)}} \rightarrow e^{-i(\pi/4)\sigma_Z^{(1)}\sigma_Z^{(2)}} e^{-i h_Z (\sigma_Z^{(1)} + \sigma_Z^{(2)})}, \quad (1.48)$$

then achieving a Bell-state fidelity above 0.99 requires $h_Z/J_{12}^{(2)} < 0.05$. For a typical coupling strength of $J_{12}^{(2)}/2\pi = 26.5 \text{ Hz}$ (realized in Ref. [163]), this translates to $h_Z/2\pi < 1.325 \text{ Hz}$.

Fortunately, since these error terms are proportional to σ_Z , several compensation strategies are available. The most straightforward approach is to absorb the systematic shift into virtual Z rotations [253], which amounts to a software update of the phase reference and incurs no physical overhead. Alternatively, in the presence of microwave driving, the detuning Δ_n in Eq. (1.37) can be adjusted to compensate the frequency shift. Finally, dynamical decoupling sequences, which are typically already employed on the MAGIC platform to protect magnetically sensitive qubit states against dephasing [239], simultaneously suppress undesired σ_Z terms without additional overhead. We discuss such noise suppression strategies via dynamical decoupling in more detail in Chapter 4.

While these techniques effectively remove the mean contribution of the unwanted local fields, fluctuations in the phonon occupation number lead to time-dependent residual field fluctuations that cannot be completely compensated in this way [205]. This underscores the importance of efficient cooling techniques, such as sympathetic cooling, to minimize phonon number fluctuations and maintain high gate fidelities, especially when scaling to longer ion chains.

Chapter 2

Higher-order effects in MAGIC-based trapped-ion platforms

In the previous chapter, we outlined the standard derivation of the effective spin–spin interaction arising in trapped-ion systems subject to a magnetic field gradient, which forms the basis of the MAGIC scheme for quantum information processing. This interaction was obtained by expanding both the external potential and the position-dependent Zeeman shift of the pseudo-spins to leading nontrivial order in the displacements from the ions’ equilibrium positions, retaining quadratic terms in the former and linear terms in the latter. While this truncation is sufficient to capture the dominant two-body Ising coupling that powers entangling operations on the MAGIC platform, any realistic device unavoidably hosts higher-order contributions whose role must be understood before scaling to large ion chains and high-fidelity applications.

Two physically distinct sources of such corrections are illustrated in Fig. 2.1. The first, shown in panel (b), originates from anharmonicities of the potential energy, in particular from the cubic term in the expansion of the Coulomb repulsion between ions, but also from possible anharmonicities in the external trapping potential itself. The second, depicted in panel (c), arises from a non-vanishing curvature of the applied magnetic field. Both mechanisms generically generate a richer set of effective interactions than the leading two-body $\sigma_Z^{(i)}\sigma_Z^{(j)}$ coupling. As we will see, these include three-body spin–spin interactions, corrections to local fields, as well as various spin–phonon and phonon–phonon couplings.

In the remainder of this chapter, we systematically investigate these higher-order contributions and estimate their magnitude using parameters representative of current MAGIC experiments. We will show that most of the resulting terms are negligible compared to the native two-body interaction, but that two contributions require careful attention: a phonon-occupation-dependent longitudinal field whose strength grows with chain length, and the well-known two-to-one phonon conversion process, both of which can be controlled by sufficient ground-state cooling and standard compensation techniques. Beyond their role as parasitic corrections, we also briefly comment on whether such higher-order terms could be purposefully engineered as a resource for quantum information processing.

All content in this chapter is based on Ref. [205].




Figure 2.1: (a) The Magnetic Gradient Induced Coupling (MAGIC) scheme for trapped-ion quantum information processing employs a magnetic field gradient, which renders the resonance frequency of the internal pseudo-spin position-dependent. This couples the spin to the spatial phonon degree of freedom, which effectively generates a spin–spin coupling that can be used for entangling gates. In this work, we analyse higher-order contributions in such a setup, in particular (b) anharmonicities in the vibrational degrees of freedom due to Coulomb repulsion and (c) curvature of the magnetic field, which can lead to additional three-body and two-body interactions, local field corrections, as well as spin–phonon couplings.

2.1 Higher-order contributions from the potential energy

This section is devoted to higher-order terms stemming from the potential energy associated with the external degrees of freedom of the ions.

We first derive the general form of the leading corrections to the Hamiltonian, giving rise to three-body spin–spin, spin–phonon, and phonon–phonon couplings (Sec. 2.1.1). Subsequently, we address specifically the effects of cubic anharmonicities in the Coulomb repulsion (Sec. 2.1.2) as well as in the external trapping potential (Sec. 2.1.4). We also discuss the generalization of these effects when considering transversal phonon modes in Sec. 2.1.3. As we show by computing their natural strength using experimentally relevant parameters, they constitute negligible perturbations or can be compensated in typical setups, though in principle a device specifically designed to enhance these higher-order contributions might leverage them for quantum information processing tasks.

2.1.1 General third-order expansion of the potential

To start with, we calculate the effect of higher-order contributions of the external potential acting on the ion crystal. In this section, we focus on the Coulomb repulsion between ions. These contributions will naturally occur in any trapped-ion setup using a phonon bus, though we discuss their effect here specifically for the MAGIC platform. We start with general derivations, which we will adapt also in the next section for anharmonicities in the trapping potential, and then specialize to the Coulomb repulsion.

Extending the expansion of the potential energy V of a 1D trapped-ion chain given in Eq. (1.21) up to third order, we obtain

$$\begin{aligned}
 V(\mathbf{z}) &= V(\mathbf{z}_0) + V^{(2)}(\mathbf{z}) + V^{(3)}(\mathbf{z}) + \mathcal{O}(\mathbf{z}^4) \\
 &= V(\mathbf{z}_0) + \sum_{i,j} \frac{A_{ij}}{2} q_i q_j + \sum_{i,j,k} \frac{B_{ijk}}{6} q_i q_j q_k + \mathcal{O}(q^4),
 \end{aligned} \tag{2.1}$$

with

$$A_{ij} = \left. \frac{\partial^2 V}{\partial q_i \partial q_j} \right|_{\mathbf{z}_0}, \tag{2.2}$$

$$B_{ijk} = \left. \frac{\partial^3 V}{\partial q_i \partial q_j \partial q_k} \right|_{\mathbf{z}_0}. \tag{2.3}$$

The quadratic terms form the harmonic part of the potential [see Eq. (1.24)], whose eigenmodes will be used to express the higher-order terms. As in the previous section, the matrix A is diagonalized by the normal modes $\mathbf{Q}$, which are related to the position operators via Eq. (1.23) and quantized through ladder operators via Eq. (1.25).

The cubic terms B_{ijk} can derive either solely from a higher-order expansion of the Coulomb potential, or additionally from anharmonicities in the trapping potential. These scenarios will be analyzed separately later. In principle, the above expansion can be continued to any desired higher order, but the resulting effects will be subleading (see also Ref. [261]).

We note that at this order of expansion of the external potential, the axial phonon modes couple to the transverse modes [229]. Here, we neglect this effect at first, as we assume (i) axial and transverse modes to typically live at different energy scales and (ii) the phonon modes to be ground-state cooled. We discuss the contributions from transverse modes separately in Section 2.1.3.

Inserting the quantized axial normal modes, the anharmonic part of Eq. (2.1) reads

$$V^{(3)} = \frac{1}{6} \sum_{i,j,k} B_{ijk} \sum_{l,r,s} S_{il} S_{jr} S_{ks} \Delta z_l \Delta z_r \Delta z_s (a_l^\dagger + a_l)(a_r^\dagger + a_r)(a_s^\dagger + a_s). \tag{2.4}$$

Adding this contribution to the Hamiltonian in Eq. (1.26) and applying the polaron transformation of Eq. (1.27) yields the total Hamiltonian $\tilde{H} = \tilde{H}^{(2)} + \tilde{H}^{(3)}$ of the system, given by

$$\tilde{H} = \tilde{H}^{(2)} + \frac{1}{6} \sum_{ijk} C_{ijk} \left(a_i^\dagger + a_i + \sum_n \epsilon_{ni} \sigma_Z^{(n)} \right) \left(a_j^\dagger + a_j + \sum_n \epsilon_{nj} \sigma_Z^{(n)} \right) \left(a_k^\dagger + a_k + \sum_n \epsilon_{nk} \sigma_Z^{(n)} \right). \tag{2.5}$$

Here, $\tilde{H}^{(2)}$ is the standard MAGIC Hamiltonian in Eq. (1.32), resulting from an expansion of the potential up to second order, and we have defined

$$C_{ijk} = \sum_{mnl} B_{mnl} S_{mi} S_{nj} S_{lk} \Delta z_i \Delta z_j \Delta z_k. \tag{2.6}$$

Sorting the terms by their order in the spin interactions, we get

$$\begin{aligned}
\tilde{H} = \tilde{H}^{(2)} + \hbar \sum_{i < j < k} J_{ijk}^{(3)} \sigma_Z^{(i)} \sigma_Z^{(j)} \sigma_Z^{(k)} \\
+ \frac{1}{2} \sum_{ijk} C_{ijk} (a_i^\dagger + a_i) \sum_{nm} \epsilon_{nj} \epsilon_{mk} \sigma_Z^{(n)} \sigma_Z^{(m)} \\
+ \frac{1}{2} \sum_{ijk} C_{ijk} (a_i^\dagger a_j^\dagger + a_i a_j + a_i^\dagger a_j + a_i a_j^\dagger) \sum_n \epsilon_{nk} \sigma_Z^{(n)} \\
+ \frac{1}{6} \sum_{ijk} C_{ijk} (a_i^\dagger a_j^\dagger a_k^\dagger + a_i a_j a_k + a_i^\dagger a_j^\dagger a_k + a_i^\dagger a_j a_k^\dagger \\
+ a_i a_j^\dagger a_k^\dagger + a_i^\dagger a_j a_k + a_i a_j^\dagger a_k + a_i^\dagger a_j a_k^\dagger). \tag{2.7}
\end{aligned}$$

One of the corrections to the standard MAGIC Hamiltonian $\tilde{H}^{(2)}$ is thus a three-body spin interaction, given by

$$\tilde{H}^{(3)} = \hbar \sum_{i < j < k} J_{ijk}^{(3)} \sigma_Z^{(i)} \sigma_Z^{(j)} \sigma_Z^{(k)} \tag{2.8}$$

with three-body interaction strength

$$\begin{aligned}
\hbar J_{ijk}^{(3)} &= \sum_{a,b,c,l,r,s} B_{abc} S_{al} S_{br} S_{cs} \Delta z_l \Delta z_r \Delta z_s \epsilon_{il} \epsilon_{jr} \epsilon_{ks} \\
&= \sum_{lrs} C_{lrs} \epsilon_{il} \epsilon_{jr} \epsilon_{ks} \\
&= \sum_{a,b,c} B_{abc} \frac{J_{ia}^{(2)}}{\partial_z \omega_a} \frac{J_{jb}^{(2)}}{\partial_z \omega_b} \frac{J_{kc}^{(2)}}{\partial_z \omega_c}. \tag{2.9}
\end{aligned}$$

In writing Eq. (2.8), we left out contributions to the local fields that arise for two coinciding indices in the sum. As we will see below, their strength, just like the three-body interaction, is negligible in realistic situations, but they could also be compensated easily through virtual Z rotations or by adjusting the frequency detuning in the presence of a microwave driving field (see Sec. 1.3.1). The other terms (second to fifth row of Eq. (2.7)) describe a coupling of phonon modes to two-body spin interactions and local fields, as well as the usual anharmonic phonon term appearing due to anharmonicities of the external potential [229].

2.1.2 Higher-order effects from the Coulomb repulsion

Up to here, the discussion applies to general anharmonicities in the external trapping potential. To obtain concrete order-of-magnitude estimates, we now specialize to the case of anharmonicities deriving solely from the Coulomb interactions, discussing each of the above additional contributions term by term. This complements the analysis of Ref. [261], where the phonon dependency was discussed within a mean field approximation assuming thermal phonon populations, with qualitatively similar conclusions.

The strength of the additional terms in Eq. (2.7) is determined by the cubic expansion

coefficients B_{ijk} . From the Coulomb repulsion, they can be evaluated as [229]

$$B_{ijk} = \frac{-6m\omega_z^2}{l} \sum_{q \neq i} \frac{u_i - u_q}{|u_i - u_q|^5} \times (\delta_{ij} - \delta_{qj})(\delta_{ik} - \delta_{qk}), \quad (2.10)$$

where l is the characteristic length scale, defined via $l^3 = e^2/(4\pi\epsilon_0 m\omega_z^2)$ [219]. We have further defined dimensionless equilibrium positions by $\mathbf{u} = \mathbf{z}_0/l$.

Table 2.1: Comparison of always-on interaction strengths in the MAGIC setup with ion chains of different length ($N = 5, 10, 15$ ions), different ion species ($^{171}\text{Yb}^+$ and $^9\text{Be}^+$), and for two different magnetic field gradients. As qubit states, we consider $|F=0, m_F=0\rangle$ and $|F=1, m_F=+1\rangle$ for $^{171}\text{Yb}^+$, and $|F=1, m_F=0\rangle$ and $|F=2, m_F=+1\rangle$ for $^9\text{Be}^+$. The table lists illustrative values of the maximal and minimal two-body coupling strength $|J_{ij}^{(2)}|$, the maximal three-body coupling strength $|J_{ijk}^{(3)}|$ due to anharmonicities in the Coulomb repulsion, as well as the largest local field correction $|h_i^z| = \sum_{jk} C_{jjk} \epsilon_{ik}$ for a ground-state cooled system (which may be amplified by higher phonon mode occupations, see Sec. 2.1.2). All interaction strengths are given in units of $2\pi \times \text{Hz}$.

(a) $\partial_z B = 19 \text{ T m}^{-1}$					
ion species	N	$\max_{ij} J_{ij}^{(2)} $	$\min_{ij} J_{ij}^{(2)} $	$\max_{ijk} J_{ijk}^{(3)} $	$\max_i h_i^z $
$^{171}\text{Yb}^+$	5	26.5	12.9	11.5×10^{-6}	6.2
$^{171}\text{Yb}^+$	10	18.6	5.8	4.2×10^{-6}	9.6
$^{171}\text{Yb}^+$	15	15.1	3.7	2.0×10^{-6}	11.9
$^9\text{Be}^+$	5	125.8	61.3	195×10^{-6}	22.2
$^9\text{Be}^+$	10	88.2	27.5	70.8×10^{-6}	34.1
$^9\text{Be}^+$	15	71.5	17.5	34.4×10^{-6}	42.5
(b) $\partial_z B = 150 \text{ T m}^{-1}$					
ion species	N	$\max_{ij} J_{ij}^{(2)} $	$\min_{ij} J_{ij}^{(2)} $	$\max_{ijk} J_{ijk}^{(3)} $	$\max_i h_i^z $
$^{171}\text{Yb}^+$	5	1650.6	805.1	0.006	49.2
$^{171}\text{Yb}^+$	10	1157.7	361.1	0.002	75.6
$^{171}\text{Yb}^+$	15	938.4	229.7	0.001	94.2
$^9\text{Be}^+$	5	7837.3	3822.9	0.1	175.1
$^9\text{Be}^+$	10	5496.8	1714.7	0.04	268.9
$^9\text{Be}^+$	15	4455.8	1090.8	0.02	335.2

Three-body spin interactions

As a concrete example, we calculate the three-body interaction strength in a one-dimensional chain of five $^{171}\text{Yb}^+$ ions in an axial harmonic trapping potential of strength $\omega_z = 2\pi \times 130 \text{ kHz}$ and an applied magnetic gradient of 150 T m^{-1} . The numerical computation is based on a microscopic model of the ion chain, where we first find the ions' equilibrium positions and then diagonalize the phonon eigenmodes (see discussion in Sec. 1.1.2), from which we

can then derive all relevant energy scales. In this example, we find the highest three-body interaction strength to be $J_{145}^{(3)} \approx 2\pi \times 0.006 \text{ Hz}$ ($J_{125}^{(3)} \approx -2\pi \times 0.006 \text{ Hz}$). Compared to the smallest two-body interaction strength in the same system, $J_{15}^{(2)} \approx 2\pi \times 805 \text{ Hz}$, the three-body spin interactions introduced by anharmonicities in the Coulomb repulsion are negligible.

Analogous numerical values for longer ion chains, as well as for a different lighter ion species ($^9\text{Be}^+$), are shown in Tab. 2.1. Even though the maximal three-body interaction strength is an order of magnitude larger for the lighter ions, it is still negligible compared to the also increased two-body interaction strengths.

Phonon-nonconserving two-body spin interactions

The term in the second row of Eq. (2.7) gives corrections of the form $\sim (a_i^\dagger + a_i)\sigma_Z^{(n)}\sigma_Z^{(m)}$ to the MAGIC two-body spin couplings. Their strength scales as $C_{ijk}\epsilon_{nj}\epsilon_{mk}$, and is thus one order in the Lamb–Dicke parameter stronger than $J_{ijk}^{(3)}$. Formally, they are of the same order in the Lamb–Dicke parameter as the leading two-body interaction $J_{ij}^{(2)}$ (Eq. (1.33)). However, the terms $\sim (a_i^\dagger + a_i)\sigma_Z^{(n)}\sigma_Z^{(m)}$ are suppressed by another factor C_{ijk} ; even more, in contrast to $J_{ij}^{(2)}$, they do not preserve phonon number, and can thus safely be neglected within a rotating wave approximation: In the interaction picture given by the diagonal phonon Hamiltonian $\hbar \sum_n \nu_n a_n^\dagger a_n$, the creation and annihilation operators gain a time dependence $a_i^\dagger \rightarrow e^{i\nu_i t} a_i^\dagger$. They thus rotate on a time scale significantly faster than $C_{ijk}\epsilon_{nj}\epsilon_{mk}$. For instance, in the above example of a chain of five $^{171}\text{Yb}^+$ ions, the strength of the terms in the second row of Eq. (2.7) is $|C_{ijk}\epsilon_{nj}\epsilon_{mk}| \lesssim 10^{-1} \text{ Hz}$. This value corresponds to interaction scales that are weaker than typical coherence times in trapped-ion hardware [262, 263]. Moreover, it is orders of magnitude smaller than even the smallest axial phonon frequency, $\omega_{\text{COM}} = \omega_z = 2\pi \times 130 \text{ kHz}$, such that this contribution can be neglected in a rotating wave approximation.

Spin-dependent phonon pair production and annihilation

The term in the third row of Eq. (2.7) has parts with different physical interpretation, which need to be discussed separately. In the same interaction picture as above, the terms creating (and annihilating) phonon pairs $a_i^\dagger a_j^\dagger$ (and $a_i a_j$) rotate with a frequency $\nu_i + \nu_j$, which is much larger than the associated energy scale $C_{ijk}\epsilon_{nk}$. Again, for the above example with five ions, we get corrections on the order of $|C_{ijk}\epsilon_{nk}| \lesssim 1 \text{ Hz}$. As the phonon frequencies are orders of magnitude larger, these terms can again safely be neglected within a rotating wave approximation.

Spin-dependent phonon hopping

In contrast, the number-preserving terms of the form $a_i^\dagger a_j$ rotate with the frequency $|\nu_i - \nu_j|$. For axial phonons and small systems, the difference of phonon mode frequencies is large (e.g., for a chain of five ions in a trapping potential of $\omega_z = 2\pi \times 130 \text{ kHz}$, phonon modes are separated by at least about $2\pi \times 80 \text{ kHz}$). Terms with $i \neq j$ thus still rotate rapidly as compared to the time scale $C_{ijk}\epsilon_{nk}$ and can be neglected.




Figure 2.2: Corrections to the local field acting on each ion n (numbered from left to right) due to the anharmonicity of the Coulomb repulsion, $\sum_{ik} C_{iik} \epsilon_{nk} \sigma_Z^{(n)}$ [third line in Eq. (2.7)]. Values computed for a chain of 15 $^{171}\text{Yb}^+$ ions and plotted over their spatial positions, with a trapping frequency of $\omega_z = 2\pi \times 130 \text{ kHz}$ and magnetic field gradients of (a) 150 T m^{-1} and (b) 19 T m^{-1} , assuming ground-state-cooled phonons. These corrections are highly inhomogeneous and can acquire significant values comparable to spin-spin couplings (see Fig. 1.1). The grey dashed line is a linear fit through the central three ions.

Phonon-occupation-dependent longitudinal field

We are thus left with terms where $i = j$, which give a correction to the local fields of strength $\sim \sum_{ik} C_{iik} \epsilon_{nk} (2n_i + 1)$, where $n_i = a_i^\dagger a_i$. Again, in the example of a chain of five $^{171}\text{Yb}^+$ ions with $\omega_z = 2\pi \times 130 \text{ kHz}$ and a magnetic gradient of 150 T m^{-1} , and assuming zero phonon number excitations ($n_i = 0$), we get corrections on the order of 10^1 Hz . These are largest for the outermost ions at the edges of the chain ($n = 1$ and $n = 5$), amounting to $\sum_{ik} C_{iik} \epsilon_{nk} \approx \pm 2\pi \times 49.2 \text{ Hz}$. For a magnetic gradient of 19 T m^{-1} , as was used in previous experiments [163], this magnitude drops to $2\pi \times 6.2 \text{ Hz}$. For non-zero phonon occupation numbers, these additional longitudinal fields get further enhanced by factors of $(2n_i + 1)$ in each term of the summation. The local field corrections are comparable to or even larger (for non-zero phonon occupation numbers) than the two-body coupling strengths $J^{(2)}/2\pi \approx 13 \text{ Hz} - 27 \text{ Hz}$. They generate an additional erroneous contribution to the two-qubit gates when implementing them as explained in Sec. 1.3.3 (the magnitude of such gate errors has also been previously discussed in Ref. [261]).

In Fig. 2.2, we show the behavior of the local field corrections across a chain of 15 ions for magnetic field gradients of (a) 150 T m^{-1} and (b) 19 T m^{-1} . The corrections are largest

towards the edges of the chain (equal magnitudes with opposite signs) and vanish for the ion at the center. The magnitude of the correction scales approximately linearly with the distance to the center of the chain, with larger deviations toward the edges. This is visualized by the grey dashed line, which is a linear fit through the centermost three ions. Comparing the largest corrections at the edges of the chain in Fig. 2.2 (b) ($\approx 2\pi \times 11.9$ Hz) to the spin–spin coupling strengths shown in Fig. 1.1 (b) ($J_{\max}^{(2)}/2\pi \approx 15.1$ Hz), we find them to be already of the same order of magnitude as the always-on spin–spin couplings for this system size.

In Fig. 2.3, we plot the scaling of these largest corrections at the edges of the chain with system size, again for the same values of the magnetic field gradient. We find that the local field corrections increase with the length of the chain. In contrast, as illustrated in Fig. 1.1, the spin–spin coupling strengths decrease. Thus, these local corrections become more relevant as the system size increases. Phenomenologically, we find the scaling behaviour to be best described by a power law with logarithmic corrections of the form $N^a \log(bN)$, with fitting parameters $a = 0.18 \pm 0.01$, $b = 1.38 \pm 0.04$ for 150 T m^{-1} and $a = 0.18 \pm 0.01$, $b = 1.35 \pm 0.07$ for 19 T m^{-1} .

For comparison, Tab. 2.1 shows the maximal (at the edges of the chain) local field corrections also for chains of lighter $^9\text{Be}^+$ ions. For a gradient of 19 T m^{-1} , the corrections are again on the same order of magnitude as the always-on two-body coupling strengths. The ratio improves with a larger magnetic field gradient, but worsens with increased chain length.

The presence of an undesired local field could constitute an issue for the precise functioning of two-qubit gates or for the correct definition of cost functions in quantum optimization [64]. Fortunately, since the associated terms are proportional to $\sigma_Z^{(n)}$, even for nonzero mean phonon excitation numbers they can be offset by adjusting the detuning of a microwave driving field from the qubit resonances or performing virtual Z gates [253] (see Sec. 1.3.1). In typical trapped-ion devices, such a compensation is often routine, even for terms many orders of magnitude stronger than in this case. For instance, in a Mølmer–Sørensen setup on the axial modes of $^{40}\text{Ca}^+$ ions, undesired off-resonant excitation of dipole transitions can induce inhomogeneous light shifts on the order of $2 \text{ kHz} - 3 \text{ kHz}$, which need to be calibrated and compensated [245]. In addition, already simple dynamical-decoupling techniques can remove undesired terms $\sim \sigma_Z^i$ [132, 261], which will be discussed in more detail in Chapter 4.

While the mean phonon number contribution to these local fields can easily be compensated, phonon number fluctuations remain problematic as they lead to fluctuating shifts. This underscores the critical importance of cooling techniques, such as sympathetic cooling [264–266], to minimize these fluctuations and ensure precise gate operations.

Phonon creation, annihilation, and conversion terms

Finally, the last two lines of Eq. (2.7) contain terms that describe the generation/annihilation of triples of phonons as well as two-to-one phonon conversion, which have been discussed before in Ref. [229]. The former terms can be neglected since they rotate with at least three times the trap frequency, which is much larger than the interaction strength C_{ijk} . The two-to-one conversion terms, however, require careful attention. As long as all modes are ground-state cooled, these terms vanish. They can also be suppressed by designing the trap parameters such that the phonon frequencies do not conspire to render such terms resonant. For the above setup with five ions and a magnetic field gradient of 150 T m^{-1} , the smallest frequency difference $|\nu_i + \nu_j - \nu_k|$ is about 27 kHz . In this specific case, these terms can thus safely be neglected even if the chain is not fully cooled into the motional ground state.




Figure 2.3: Scaling of the largest correction $\sum_{ik} C_{iik}\epsilon_{1k}\sigma_Z^{(1)}$ (for the ion at the leftmost edge of the chain) to the local fields with increasing number of $^{171}\text{Yb}^+$ ions in the chain. Calculations were done for a vanishing phonon occupation, a trapping frequency of $\omega_z = 2\pi \times 130 \text{ kHz}$, and magnetic gradients of (a) $150 \frac{\text{T}}{\text{m}}$ and (b) $19 \frac{\text{T}}{\text{m}}$. The absolute strength increases with chain length, rendering this correction significant in particular when scaling to larger ion chains. Numerical fits (grey dashed lines) show a scaling $\sim N^{0.18} \log(1.38N)$.

Nevertheless, these caveats need to be kept in mind especially when scaling up to large ion chains, where ground-state cooling becomes more challenging and resonances can appear more easily.

To summarize this part, among the various corrections to the MAGIC trapped-ion scheme that derive from a third-order expansion of the Coulomb repulsion, only phonon-number-dependent longitudinal fields as well as two-to-one phonon conversions can become sizable. The effect of the former can be compensated via additional virtual Z rotations in the case of discrete quantum gates [253] or, in applications involving microwave driving, by adjusting the frequency detuning, while the latter can be avoided by ground-state cooling of the phonon modes.

2.1.3 Influence of transversal phonon modes

In the previous subsection, we restricted the analysis of cubic Coulomb anharmonicities to the purely axial modes. As discussed in Sec. 1.2.1, the axial and transversal modes decouple at second order in the expansion of the potential, but this is no longer the case once cubic terms are retained. Here, we briefly examine the resulting axial–transversal couplings and show that they do not alter the conclusions of the preceding analysis.

CHAPTER 2. HIGHER-ORDER EFFECTS IN MAGIC-BASED TRAPPED-ION PLATFORMS

For a one-dimensional ion chain in a harmonic trap, the third-order Coulomb contribution to the potential energy reads [229]

$$V^{(3)}(\mathbf{q}) = \sum_{i,j,k} \tilde{B}_{ijk} q_k^{(3)} \left(2 q_i^{(3)} q_j^{(3)} - 3 q_i^{(2)} q_j^{(2)} - 3 q_i^{(1)} q_j^{(1)} \right), \quad (2.11)$$

with

$$\tilde{B}_{ijk} = \frac{m\omega_z^2}{2l} \sum_{q \neq i} \frac{\text{sgn}(u_q - u_i) (\delta_{ij} - \delta_{qj})(\delta_{ik} - \delta_{qk})}{(u_q - u_i)^4}, \quad (2.12)$$

where $l^3 = e^2/4\pi\epsilon_0 m\omega_z^2$ is the characteristic length scale and $\mathbf{u} = \mathbf{q}_0/l$ are the dimensionless equilibrium positions. The first term in Eq. (2.11) involves only axial displacements and reproduces the contributions analyzed in Sec. 2.1.2. The remaining two terms couple a single axial displacement to a pair of transversal displacements in the same transversal direction, reflecting the fact that small radial excursions of the ions modify the Coulomb potential energy only at second order in the transversal coordinates.

Quantizing the modes and applying the polaron transformation as in Sec. 2.1.2, the axial-transversal contributions generate two qualitatively new types of terms. The first are spin-dependent two-phonon processes in the transversal modes,

$$\sum_{\alpha=1}^2 \sum_{ijk} C_{ijk}^{\alpha} (a_{\alpha,j}^{\dagger} a_{\alpha,k}^{\dagger} + a_{\alpha,j} a_{\alpha,k} + a_{\alpha,j}^{\dagger} a_{\alpha,k} + a_{\alpha,j} a_{\alpha,k}^{\dagger}) \sum_n \epsilon_{ni}^{(3)} \sigma_Z^{(n)}, \quad (2.13)$$

with coefficients

$$C_{ijk}^{\alpha} = -3 \Delta z_i \Delta \alpha_j \Delta \alpha_k \sum_{mnp} \tilde{B}_{mnp} S_{pi}^{(3)} S_{mj}^{(\alpha)} S_{nk}^{(\alpha)}, \quad (2.14)$$

where $S^{(\alpha)}$ is the mode participation matrix for direction α and $\Delta \alpha_j = \sqrt{\hbar/2m\nu_j^{(\alpha)}}$. As in Sec. 2.1.2, the phonon-non-conserving contributions can be dropped within a rotating-wave approximation, leaving a correction to the local fields of the form

$$\sum_n \sigma_Z^{(n)} \sum_{\alpha=1}^2 \sum_{ij} C_{ijj}^{\alpha} \epsilon_{ni}^{(3)} (2a_{\alpha,j}^{\dagger} a_{\alpha,j} + 1). \quad (2.15)$$

This contribution is the transversal analogue of the phonon-occupation-dependent longitudinal field identified in Sec. 2.1.2, but it is suppressed relative to the purely axial correction by a factor $\nu^{(3)}/\nu^{(\alpha)}$, since the transversal trap frequencies are typically much larger than the axial one [163, 245]. Its mean effect can again be absorbed by adjusting the microwave detuning or via virtual Z rotations, while residual fluctuations of the transversal phonon occupation can still degrade gate fidelity unless the transversal modes are sufficiently cooled.

The second new contribution is a three-body axial-transversal mode coupling,

$$\sum_{\alpha=1}^2 \sum_{ijk} C_{ijk}^{\alpha} (a_{3,i}^{\dagger} + a_{3,i})(a_{\alpha,j}^{\dagger} + a_{\alpha,j})(a_{\alpha,k}^{\dagger} + a_{\alpha,k}), \quad (2.16)$$

which describes the conversion between axial and transversal phonon excitations. For generic trap parameters and chain lengths, the rotating-wave approximation removes these terms as

well, but specific combinations of axial and transversal mode frequencies can render some of them resonant [229]. In such cases, the corresponding processes can still be rendered harmless by cooling the relevant modes close to their motional ground state.

In summary, the inclusion of transversal modes does not introduce qualitatively new error mechanisms beyond those identified in the purely axial analysis. The dominant transversal correction is again a phonon-occupation-dependent longitudinal field, parametrically smaller than its axial counterpart, and the resonant axial–transversal conversion processes can be suppressed by ground-state cooling. The conclusions of Sec. 2.1.2 therefore remain valid in the full three-dimensional setting.

2.1.4 Cubic anharmonicities of the external trapping potential

In addition to the intrinsic anharmonicities given by the Coulomb potential, there could also be anharmonicities stemming from the external trapping potential. While such effects are often undesired for quantum logic, anharmonicities can nowadays also be designed and exploited in a controlled way [267, 268]. Such anharmonic potentials can change the equilibrium positions of the ions and therefore modify the two-body interactions induced by a magnetic field gradient. This effect has been investigated in Ref. [223]. Here, we discuss the role of higher-order terms generated by anharmonicities in the external potential.

To this end, we consider a cubic contribution to the potential energy of the form

$$V^{(3)}(\mathbf{z}) = \sum_n \frac{\alpha_n}{6} (z_n - z_{n,0})^3, \quad (2.17)$$

where α_n describes the strength of the cubic anharmonicity felt by ion n . Such terms in the potential energy generate additional three-body spin–spin, spin–phonon, and phonon–phonon interactions, for which the general considerations discussed in Sec. 2.1.2 apply. Here, we investigate the resulting three-body spin interaction in more detail.

Inserting the expansion coefficients $B_{ijk} = \delta_{ij}\delta_{jk}\alpha_i$ into Eq. (2.9), the three-spin coupling coefficients specialize to

$$\hbar J_{ijk}^{(3)} = \sum_n \frac{\alpha_n}{(\partial_z \omega_n)^3} J_{in}^{(2)} J_{jn}^{(2)} J_{kn}^{(2)}. \quad (2.18)$$

In order to obtain a simple order-of-magnitude estimate for their typical strength, we first estimate the typical strength of the two-body spin interaction in Eq. (1.33). For this purpose, we consider only the contribution from the center-of-mass (cm) mode, whose frequency coincides with the trap frequency, $\nu_{\text{cm}} = \omega_z$, and whose matrix elements $S_{i,\text{cm}}$ are given by $1/\sqrt{N}$. Assuming that the gradient of the resonance frequency is the same for all ions ($\partial_z \omega_i = \partial_z \omega$), we have $J^{(2)} \simeq \omega_z \epsilon^2$ with the effective Lamb–Dicke parameter $\epsilon = \partial_z \omega \Delta z / \omega_z \sqrt{N}$ and oscillator length $\Delta z = \sqrt{\hbar/2m\omega_z}$. Combining this expression with Eq. (2.18), we can estimate the typical strength of the three-body coupling as

$$\hbar J^{(3)} \simeq \frac{1}{\sqrt{N}} \frac{\alpha \Delta z^3}{\hbar \omega_z} \epsilon \hbar J^{(2)} = \frac{1}{\sqrt{N}} \alpha \Delta z^3 \epsilon^3, \quad (2.19)$$

where we have made the additional simplifying assumption that the anharmonicity is the same at each ion position ($\alpha_n = \alpha$).

One may be tempted to engineer cubic terms in the potential similar to Eq. (2.17) with the goal of harnessing the resulting three-body spin interaction for quantum information processing tasks. Unfortunately, as we argue in what follows, this endeavor is unlikely to be advantageous in practice.

To this end, we note that cubic contributions to the potential energy may severely affect the stability of the trap. For instance, an external potential of the form $V_{\text{trap}}(\mathbf{z}) = \sum_n m\omega_z^2 z_n^2/2 + \sum_n \alpha z_n^3/6$ with global anharmonicity α gives rise to terms as in Eq. (2.17) with $\alpha_n = \alpha$. However, such a potential also shifts the equilibrium positions of the ions. Thus, in order to prevent ions from being ejected from the trap, the anharmonicity should be a somewhat small perturbation on top of the harmonic part, such that the potential still features a local minimum that is sufficiently deep and wide to trap the ions. That is, the energy scale $\alpha\Delta z^3$ should be well below the harmonic trapping energy $\hbar\omega_z$, requiring $\alpha\Delta z^3/\hbar\omega_z \ll 1$. On the other hand, the three-body interaction should be sufficiently strong in order to be of practical use. For example, a native three-qubit gate should outperform the equivalent gate synthesized from a combination of two-qubit gates. Thus, it is desirable that $J^{(3)} = \lambda J^{(2)}$ with λ as large as possible (though typically well below unity). The estimate in Eq. (2.19) can then be expressed as $\alpha\Delta z^3/\hbar\omega_z = \lambda\sqrt{N}/\epsilon$. As noted before, stability of the trapping potential requires this quantity to be sufficiently small, enforcing the hierarchy $\lambda = J^{(3)}/J^{(2)} \ll \epsilon/\sqrt{N}$. This means that the three-qubit interaction is bound to be weaker than the two-qubit interaction by a factor well below $\epsilon/\sqrt{N}$, which in realistic setups amounts to a suppression by at least one, if not several, orders of magnitude.

To circumvent the stability issues described above, one may attempt to engineer the terms in Eq. (2.17) *locally*, i.e., without affecting the equilibrium positions of the ions, as may be possible, e.g., in segmented ion traps [267]. However, from a practical point of view, the electrodes creating such local potentials cannot be placed arbitrarily close to the ions, but require a certain minimum distance, which is typically large compared to the spacing between ions [267]. Consequently, it becomes difficult to create localized features varying on the scale of the oscillator length Δz , which is much smaller than the distance between ions. For example, it appears challenging to engineer localized external potentials stronger than the intrinsic Coulomb potential created by the ions themselves. Even if the strength of the local cubic potential was comparable to the characteristic energy scale of the anharmonic part of the Coulomb repulsion in Eq. (2.10), $\alpha\Delta z^3 = 6m\omega_z^2\Delta z^3/l \approx 0.004\hbar\omega_z$, the maximum three-body coupling strength would amount to only $J_{\text{max}}^{(3)}/2\pi \approx 0.3\text{ Hz}$ in a chain of $N = 5$ ions with trapping frequency $\omega_z/2\pi = 130\text{ kHz}$ and magnetic field gradient $\partial_z B = 150\text{ T m}^{-1}$, which is too low for typical applications. For comparison, a potential with a *global* cubic anharmonicity of the same strength would induce equally weak three-body interactions, but would not be able to trap the ions for the above parameters in its local minimum.

To sum up, while cubic anharmonicities can in principle be a route to engineering three-body interactions in the MAGIC setup, the comparatively slow three-qubit interaction rate and the substantial experimental overhead of realizing the required cubic anharmonicities in a stable way limit the practical use of this scheme. After all, our discussion reveals that moderate cubic anharmonicities typically have a negligible parasitic impact on the standard MAGIC setup.

2.2 Contributions from magnetic field curvature

While anharmonicities in the external potential of the ion crystal appear in any trapped-ion platform, a higher-order contribution specific to the MAGIC setup originates from higher-order spatial derivatives of the magnetic field. To evaluate their effect, we expand the

magnetic-field-dependent resonance frequency of the ions [see Eq. (1.19)] up to second order,

$$\omega(z_n) = \omega_n + \partial_z \omega(z_n) \Big|_{z_{n,0}} (z_n - z_{n,0}) + \frac{1}{2} \partial_z^2 \omega(z_n) \Big|_{z_{n,0}} (z_n - z_{n,0})^2 + \mathcal{O}(z_n^3). \quad (2.20)$$

The Hamiltonian H_{int} in Eq. (1.20), describing the internal energy of the ion chain, then becomes

$$H_{\text{int}} = -\frac{\hbar}{2} \sum_n \left(\omega_n + \partial_z \omega_n q_n + \frac{1}{2} \partial_z^2 \omega_n q_n^2 \right) \sigma_Z^{(n)}. \quad (2.21)$$

Inserting again the quantized normal modes, the internal energy Hamiltonian reads

$$\begin{aligned} H_{\text{int}} = & -\frac{\hbar}{2} \sum_n \omega_n \sigma_Z^{(n)} - \frac{\hbar}{2} \sum_{nl} \partial_z \omega_n S_{nl} \Delta z_l (a_l^\dagger + a_l) \sigma_Z^{(n)} \\ & - \frac{\hbar}{4} \sum_{nlr} \partial_z^2 \omega_n S_{nl} S_{nr} \Delta z_l \Delta z_r (a_l^\dagger + a_l) (a_r^\dagger + a_r) \sigma_Z^{(n)}. \end{aligned} \quad (2.22)$$

Performing, as before, the polaron transformation to decouple phonons and spins in the first order terms, Eq. (1.27), we obtain the total Hamiltonian $\tilde{H} = \tilde{H}^{(2)} + \tilde{H}^{(3)}$ with

$$\tilde{H}^{(3)} = -\frac{\hbar}{4} \sum_{nlr} \partial_z^2 \omega_n S_{nl} S_{nr} \Delta z_l \Delta z_r \sigma_Z^{(n)} \left(a_l^\dagger + a_l + \sum_i \epsilon_{il} \sigma_Z^{(i)} \right) \left(a_r^\dagger + a_r + \sum_j \epsilon_{jr} \sigma_Z^{(j)} \right). \quad (2.23)$$

Again, $\tilde{H}^{(2)}$ is the standard MAGIC Hamiltonian, while $\tilde{H}^{(3)}$ are the corrections due to the expansion of the magnetic field up to third order. We can sort all terms by their power in the spin operators and write

$$\begin{aligned} \tilde{H} = & \tilde{H}^{(2)} - \frac{\hbar}{4} \sum_{ijk} J_{ijk}^{(3)} \sigma_Z^{(i)} \sigma_Z^{(j)} \sigma_Z^{(k)} \\ & - \frac{\hbar}{2} \sum_{ijlr} \tilde{C}_{ilr} \epsilon_{jr} (a_l^\dagger + a_l) \sigma_Z^{(i)} \sigma_Z^{(j)} \\ & - \frac{\hbar}{4} \sum_{ilr} \tilde{C}_{ilr} (a_l^\dagger a_r^\dagger + a_l^\dagger a_r + a_l a_r^\dagger + a_l a_r) \sigma_Z^{(i)}, \end{aligned} \quad (2.24)$$

where we defined

$$\tilde{C}_{ilr} = \partial_z^2 \omega_i S_{il} S_{ir} \Delta z_l \Delta z_r. \quad (2.25)$$

Analogously to the discussion of Sec. 2.1.2, the terms in the second row give corrections to the two-body spin interactions of the form $(a_l^\dagger + a_l) \sigma_Z^{(i)} \sigma_Z^{(j)}$. Again, these do not preserve phonon numbers and can be neglected within a rotating wave approximation. The terms in the third row only give corrections to local fields of strength $\sim \tilde{C}_{irr} (2n_r + 1)$ with $n_r = a_r^\dagger a_r$ (neglecting phonon-number-non-preserving terms as before), which can be offset by microwave detunings or neglected for small magnetic field curvatures ($\tilde{C}_{irr} \propto \partial_z^2 \omega_i$). Due to the different origin of the higher-order terms as compared to Sec. 2.1.2, there are no cubic phonon terms in this expansion.

The remaining effective three-spin interaction is

$$\tilde{H}^{(3)} = -\frac{\hbar}{4} \sum_{ijk} J_{ijk}^{(3)} \sigma_Z^{(i)} \sigma_Z^{(j)} \sigma_Z^{(k)} \quad (2.26)$$

with the three-body coupling strength

$$J_{ijk}^{(3)} = \sum_{lr} \partial_z^2 \omega_i S_{il} S_{ir} \Delta z_l \Delta z_r \epsilon_{jl} \epsilon_{kr}. \quad (2.27)$$

The form of this term differs from the ones in the previous section [Eqs. (2.9) and (2.18)], since it did not derive from a contribution $\propto x^3$, but rather $\propto x^2 \sigma^z$.

By defining $\gamma_n = \partial_z^2 \omega_n / (\partial_z \omega_n)^2$, i.e., the ratio between the magnetic field curvature and the square of the gradient, and using the spin–spin coupling strength given in Eq. (1.33), one can write this interaction even more compactly as

$$J_{nij}^{(3)} = \gamma_n J_{ni}^{(2)} J_{nj}^{(2)}. \quad (2.28)$$

Notably, $J_{nij}^{(3)} = J_{nji}^{(3)}$, but in general $J_{nji}^{(3)} \neq J_{jni}^{(3)}$. One can symmetrize the 3-body coupling strength to

$$\tilde{J}_{ijk}^{(3)} = \gamma_i J_{ij}^{(2)} J_{ik}^{(2)} + \gamma_j J_{ji}^{(2)} J_{jk}^{(2)} + \gamma_k J_{ki}^{(2)} J_{kj}^{(2)}, \quad (2.29)$$

which assumes a form similar to that obtained in Ref. [269] for a trapped-ion system driven close to the second sideband.

In analogy to the procedure for anharmonic external potentials in Sec. 2.1.4, the typical strength of the three-body coupling induced by the magnetic field curvature can be estimated as

$$J^{(3)} = \frac{1}{\sqrt{N}} \frac{\partial_z^2 \omega \Delta z}{\partial_z \omega} \epsilon J^{(2)} = \frac{1}{N} \partial_z^2 \omega \Delta z^2 \epsilon^2. \quad (2.30)$$

Remarkably, this quantity scales quadratically with the effective Lamb–Dicke parameter ϵ and thus with the magnetic field gradient, similar to the two-body coupling $J^{(2)}$.

For a potential application of this “magnetic curvature induced three-body coupling” in quantum information processing, the second derivative of the local resonance frequency $\partial_z^2 \omega$ over length scales on the order Δz should be sufficiently large. However, similarly as for the local cubic potentials discussed in Sec. 2.1.4, in practice it is very hard to engineer magnetic field distributions with sufficiently sharp localized features.

Instead, we consider in what follows a *global* magnetic field curvature of the form $\mathbf{B}(z) = \mathbf{B}_0 + (\partial_z B z + \partial_z^2 B z^2/2) \hat{\mathbf{e}}_z + \mathcal{O}(z^3)$. Unlike higher-order terms in the external potential, higher-order contributions to the magnetic field do not alter the equilibrium positions of the ions as they act on the internal degrees of freedom, but they can lead to a position-dependent magnetic-field gradient $\partial_z B(z_n)$. As an example, we consider a magnetic field curvature of $\partial_z^2 B \simeq 10^3 \text{ T m}^{-2}$, which can be reached using suitably designed Halbach magnets [270]. If the ion chain is centered around the origin, the variation of the magnetic field gradient is negligible, $\partial B(z_n)/\partial z = \partial_z B + z_n \partial_z^2 B \approx \partial_z B$, since the spacing between ions is on the order of $10 \mu\text{m}$. In addition, for five $^{171}\text{Yb}^+$ ions with a magnetic field gradient of $\partial_z B = 150 \text{ T m}^{-1}$ and a trap frequency of $\omega_z/2\pi = 130 \text{ kHz}$, we obtain three-body coupling strengths of only $J^{(3)} \sim 2\pi \times 10^{-5} \text{ Hz}$, which is orders of magnitude too low for practical use. Nonetheless, this example underlines that, as long as the magnetic field gradient does not significantly vary on length scales comparable to the oscillator length Δz , parasitic effects from nonlinear gradients can safely be neglected in the standard MAGIC scheme.

2.3 Comparison to existing schemes for generating higher-order interactions on trapped-ion platforms

As shown above, higher-order terms in the MAGIC setup naturally lead to three-body interactions and spin-phonon couplings. We have found that the three-body interaction strength is negligible compared to the two-body interactions for realistic experimental setups. For instance, the corrections arising due to inherent anharmonicities in the Coulomb repulsion only produce a three-body interaction strength on the order of 10^{-2} Hz for typical parameters (see Sec. 2.1.2).

However, since such types of interactions are key ingredients in applications ranging from quantum simulation of topological and Peierls phase transitions [271, 272] and lattice gauge theories [273–275] to improved schemes of quantum annealing [203, 276], one may wonder if there are alternative feasible ways to engineer such interactions and how the achievable interaction strengths compare to the ones we have discussed above.

Indeed, various proposals have been designed for this purpose specifically with trapped-ion setups in mind. One rather generally applicable approach relies on obtaining three-body interactions perturbatively from two-body terms, e.g., via an energy penalty [277] or via a high-frequency expansion based on Floquet engineering [278]. The 3-body interaction in Ref. [277] arises from perturbation theory, combining an effective magnetic field acting on the pseudo-spins, $\sim K$, and a Mølmer-Sørensen interaction $\sim J$. These interactions are further suppressed by another Mølmer-Sørensen interaction $\sim V$, generated by an additional pair of beams, which is assumed to be the largest energy scale. Within perturbation theory, the 3-body strength is therefore of order KJ/V . The numerics in Ref. [277] show that values of around $J = K = 0.3 V$ are sufficient to enter the relevant perturbative regime. Assuming that V corresponds to reasonably achievable interaction strengths from Mølmer-Sørensen gates on the order of $\sim 2\pi \times 1$ kHz, this results in a three-body interaction strength on the order of $0.3^2 V \approx 2\pi \times 0.1$ kHz. As a feature of the approach in Ref. [277], in the effective Hamiltonian the two-body interactions play no role any more.

Other approaches are based on continuously and simultaneously addressing the first and second sideband [269, 274], which results in the simultaneous presence of terms that are formally similar to the terms $q_n \sigma_Z^{(n)}$ and $q_n^2 \sigma_Z^{(n)}$ that appear in Eq. (2.21). In this approach, the ability to use large Rabi frequencies on the first and second sideband enables large three-body interactions whose strength can be tuned to match that of the two-body interactions. The authors of Ref. [269] discuss how the dipole force can be adapted, assuming reasonable trap frequencies ~ 10 MHz, laser detunings ~ 1.25 MHz, and stiffness parameter ~ 0.05 , to reach a three-body interaction strength of about $2\pi \times 0.1$ kHz, and, by relaxing trap frequencies or designing specific microtraps, even up to about $2\pi \times 1.6$ kHz [269].

A further promising proposal has recently been put forward based on a combination of discrete, spin-dependent displacement and squeezing pulses [279]. In a chain of four $^{171}\text{Yb}^+$ ions, the authors estimate the gate time of a fully entangling three-qubit gate, $U = \exp(-i \frac{\pi}{4} \sigma_Z^{(i)} \sigma_Z^{(j)} \sigma_Z^{(k)})$, to be $T \approx 130 \mu\text{s}$ for feasible setups with laser-based control, corresponding to a three-body coupling strength on the order of $2\pi \times 1$ kHz.

The above three proposals rely only on ingredients already demonstrated in trapped ions, though their combination to obtain a three-body interaction has thus far not been demonstrated in experiment.

Finally, a standard way that is always possible given a universal gate set is to engineer unitary three-qubit interactions through a combination of two-qubit and single-qubit gates. For example, a term of the form $\sigma_Z^{(i)} \sigma_Z^{(j)} \sigma_Z^{(n)}$ can be realized using four Mølmer-Sørensen

gates; an explicit gate-level construction is given in Sec. 5.1.5. Neglecting the gate time of single-qubit operations, one can associate an effective three-body strength to this gate sequence of order $J^{(3)} = J^{(2)}/4$, which is a rather sizable interaction strength. As the unitary three-body gate is generated exactly in this case, there are no additional two-body interactions.

In addition to three-spin interactions, there are also interesting applications involving higher-order products of spin and phonon operators, such as terms of the form $a_i a_j \sigma_Z^{(n)}$ or $a_i \sigma_Z^{(n)} \sigma_Z^{(m)}$. Above, we have seen that the higher-order expansion can generate such terms as small corrections. In Ref. [273], it has been shown how these terms can be used to quantum simulate a lattice version of quantum electrodynamics. In that work, it was found that near-resonant addressing of first and second sidebands may generate interaction strengths up to $2\pi \times 120$ Hz.

Chapter 3

Quantum optimization

This chapter establishes the theoretical background on quantum annealing and related optimization algorithms used in the remainder of this thesis, and introduces the specific optimization problems that will serve as benchmarks in later chapters. The first part, Sec. 3.1, develops the framework of quantum annealing, covering the encoding of classical optimization problems into cost Hamiltonians, the standard annealing protocols, the adiabatic theorem, and the role of the relevant energy scales in determining annealing performance. The second part, Sec. 3.2, presents a collection of combinatorial optimization problems that will be revisited throughout the thesis. These range from the analytically tractable paradigmatic ferromagnetic p -spin model and a geometrically local Ising model, over problems with direct industrial relevance such as multiple object tracking and the cutting stock problem, to the canonical NP-complete 3-SAT problem and a selection of further problems whose natural formulation is polynomial rather than quadratic. The chapter closes with Sec. 3.3, in which we briefly review two gate-based algorithms closely related to quantum annealing, Trotterized quantum annealing and the quantum approximate optimization algorithm (QAOA), which provide the digital counterparts to the analog protocols discussed earlier.

This chapter is based on material in Refs. [144, 203, 204].

3.1 Quantum annealing

In this section, we introduce the basic framework of quantum annealing as used throughout this thesis. We first describe in Sec. 3.1.1 how a classical combinatorial optimization problem is encoded into a quantum cost Hamiltonian whose ground state represents the solution. We then introduce in Sec. 3.1.2 the standard annealing protocols considered in the rest of the work, and discuss in Sec. 3.1.3 the adiabatic theorem, which sets a sufficient condition on the sweep duration for the system to remain in its instantaneous ground state during the evolution. The remaining subsections analyze the characteristic energy scale of the cost Hamiltonian (Sec. 3.1.4) and how to choose the strength of the driving Hamiltonian relative to it (Sec. 3.1.5).

We briefly remark on the terminology used in this thesis. Adiabatic quantum computation (AQC), in its original formulation [30], refers to a model of quantum computation in which the system is initialized in the ground state of a driving Hamiltonian and slowly evolved toward the ground state of a problem Hamiltonian, with the adiabatic theorem guaranteeing that excited states remain unpopulated throughout the protocol. We use the term quantum annealing in a broader sense: it includes AQC as the strict-adiabatic limit,

but also encompasses protocols in which transitions to higher excited states are allowed, whether as a deliberate computational resource or as an unavoidable consequence of finite sweep durations, environmental noise, and limited control fidelity on physical hardware. Throughout this thesis, we use the two terms interchangeably whenever the distinction is unnecessary, and reserve the adjective “adiabatic” for statements that explicitly assume or require strict ground-state evolution.

3.1.1 From classical cost functions to cost Hamiltonians

The starting point of quantum annealing is a classical combinatorial optimization problem formulated over a set of N binary variables $x_i \in \{0, 1\}$, $i = 1, \dots, N$. The objective is to minimize a real-valued cost function $C(\mathbf{x})$ defined on the Boolean hypercube $\{0, 1\}^N$. When $C(\mathbf{x})$ is at most quadratic in the binary variables, the problem is known as a quadratic unconstrained binary optimization (QUBO) problem and can be written in the canonical form

$$C(\mathbf{x}) = \sum_{i \leq j} Q_{ij} x_i x_j, \quad (3.1)$$

with real coefficients Q_{ij} . Diagonal entries Q_{ii} encode local biases (since $x_i^2 = x_i$ for binary variables), while off-diagonal entries encode pairwise couplings between variables. More generally, when $C(\mathbf{x})$ contains products of three or more binary variables, the problem is referred to as a polynomial unconstrained binary optimization (PUBO) problem, with k -body terms of the form $Q_{i_1 \dots i_k}^{(k)} x_{i_1} \dots x_{i_k}$. Cost functions of PUBO type are discussed in Chapter 5.

To make the optimization problem accessible to the quantum annealing algorithm, the classical binary variables are promoted to mutually commuting quantum-mechanical operators acting on a register of N qubits. Throughout the majority of this thesis we adopt the Pauli convention, in which each binary variable is identified with a Pauli- z operator $\hat{\sigma}_i^z$ (with eigenvalues ± 1) via

$$x_i \longrightarrow \frac{1 - \hat{\sigma}_i^z}{2}, \quad (3.2)$$

so that the eigenvalue $+1$ of $\hat{\sigma}_i^z$ corresponds to $x_i = 0$ and the eigenvalue -1 to $x_i = 1$. An equivalent alternative used in Chapter 5 is the spin- $\frac{1}{2}$ convention, in which the binary variables are first mapped to classical Ising spins $s_i = \frac{1}{2} - x_i \in \{\pm \frac{1}{2}\}$ and then promoted to spin operators $\hat{s}_i^z$ with eigenvalues $\pm \frac{1}{2}$,

$$x_i \longrightarrow \frac{1}{2} - \hat{s}_i^z. \quad (3.3)$$

The two mappings are related by $\hat{s}_i^z = \hat{\sigma}_i^z / 2$ and are therefore equivalent up to a rescaling of the resulting coupling constants; in particular, neither mapping changes the degree of the underlying optimization problem.

Substituting Eq. (3.2) into the QUBO cost function (3.1) and using $(\hat{\sigma}_i^z)^2 = 1$ yields the cost Hamiltonian

$$\hat{H}_{\text{cost}} = \sum_{i < j} J_{ij} \hat{\sigma}_i^z \hat{\sigma}_j^z + \sum_i h_i^z \hat{\sigma}_i^z, \quad (3.4)$$

which is diagonal in the computational basis. The effective two-body couplings J_{ij} and longitudinal fields h_i^z are linear combinations of the original QUBO coefficients,

$$J_{ij} = \frac{Q_{ij}}{4}, \quad h_i^z = -\frac{Q_{ii}}{2} - \frac{1}{4} \sum_{j \neq i} Q_{ij}, \quad (3.5)$$

where in the second expression Q_{ij} is understood as symmetrized, $Q_{ij} = Q_{ji}$ for $i \neq j$. The substitution also produces an overall constant, which we omit since it merely shifts the spectrum and does not affect the eigenstates of $\hat{H}_{\text{cost}}$. By construction, the ground state of $\hat{H}_{\text{cost}}$ encodes the solution of the original QUBO problem. Higher-order cost functions of PUBO type translate analogously into cost Hamiltonians containing k -body products of $\hat{\sigma}^z$ operators of the same degree as the original polynomial, as we discuss in Sec. 3.2.5 for the 3-SAT problem and in Chapter 5.

In certain situations, it is advantageous to encode the optimization problem entirely in terms of two-body (and higher-order) interactions by absorbing the local fields into couplings with an additional ancilla qubit a . This is achieved by the substitution $h_i^z \hat{\sigma}_i^z \rightarrow J_{ai} \hat{\sigma}_a^z \hat{\sigma}_i^z$ with $J_{ai} = h_i^z$, yielding the modified cost Hamiltonian

$$\hat{H}_{\text{cost}} = \sum_{i < j} J_{ij} \hat{\sigma}_i^z \hat{\sigma}_j^z + \hat{\sigma}_a^z \sum_i J_{ai} \hat{\sigma}_i^z. \quad (3.6)$$

The driving Hamiltonian is correspondingly extended to act on the ancilla as well, and the system is initialized in the equal superposition over all $N + 1$ qubits. As a consequence of this reformulation, $\hat{H}_{\text{cost}}$ is invariant under a global flip of all $\hat{\sigma}^z$ operators (including the ancilla), so that a previously unique ground state of the original cost Hamiltonian becomes two-fold degenerate. Either of the two ground states yields the solution to the original problem, since the measured ancilla outcome determines whether the remaining qubits must be flipped. A key benefit of this encoding is that fluctuating local fields, which constitute a significant noise source on, for example, trapped-ion platforms (see Chapter 1), are decoupled from the problem itself and can be mitigated by dynamical decoupling sequences applied during the annealing protocol, as discussed in Chapter 4.

3.1.2 Quantum annealing protocols

The basic idea of quantum annealing is to evolve the system from an easily prepared ground state of a driving Hamiltonian $\hat{H}_{\text{drive}}$ into the desired ground state of $\hat{H}_{\text{cost}}$, which encodes the solution to a classical combinatorial optimization problem (see previous section), by adiabatically changing Hamiltonian parameters. The driving Hamiltonian must not commute with the cost Hamiltonian, and its ground state should be easy to prepare. A standard choice is a uniform transverse field,

$$\hat{H}_{\text{drive}} = -h^x \sum_i \hat{\sigma}_i^x, \quad (3.7)$$

whose ground state is the uniform superposition $|+\rangle^{\otimes N}$ (for $h^x > 0$). Below we discuss two simple annealing protocols which will be utilized in the rest of this work. Section 3.1.3 then explains the adiabatic theorem, which gives a bound on how fast a quantum annealing algorithm can be run.

The standard annealing protocol considers a linear sweep from $\hat{H}_{\text{drive}}$ to $\hat{H}_{\text{cost}}$, where the instantaneous Hamiltonian reads

$$\hat{H}(s) = (1 - s) \hat{H}_{\text{drive}} + s \hat{H}_{\text{cost}}, \quad (3.8)$$

where $s = t/T \in [0, 1]$ is the time in units of the sweep duration T . At $s = 0$, the Hamiltonian reduces to the drive Hamiltonian and the system is initialized in its ground state $|+\rangle^{\otimes N}$. By the adiabatic theorem (see next section), if the sweep is sufficiently slow, the system

remains in the instantaneous ground state throughout the evolution and reaches the ground state of $\hat{H}_{\text{cost}}$ at $s = 1$.

An alternative formulation keeps the cost Hamiltonian at unit strength throughout the annealing protocol and only tunes down the strength of the drive Hamiltonian:

$$\hat{H}(s) = \hat{H}_{\text{cost}} + C(s) \hat{H}_{\text{drive}}, \quad (3.9)$$

where $C(s)$ is a monotonically decreasing control function satisfying $C(0) = 1$ and $C(1) = 0$. For the standard linear schedule, $C(s) = 1 - s$. Unlike the symmetric interpolation (3.8), the cost Hamiltonian is present from the outset, so the initial ground state of $\hat{H}(0) = \hat{H}_{\text{cost}} + C(0) \hat{H}_{\text{drive}}$ is no longer exactly the uniform superposition $|+\rangle^{\otimes N}$. Instead, $|+\rangle^{\otimes N}$ is only an approximate ground state whose accuracy improves with increasing strength of the driver Hamiltonian h_x ; in the limit $h_x \rightarrow \infty$, the driver dominates and the exact ground state is recovered. In practice, h_x is chosen large enough so that the resulting infidelity of the initial state is negligible compared to other sources of error. At the end of the anneal ($s = 1$), the Hamiltonian reduces to the cost Hamiltonian and the solution state is reached if the speed of the protocol satisfies the adiabatic theorem. This modification of the standard quantum annealing protocol can be useful in situations where the strength of the cost Hamiltonian can not be modified, or it is more convenient to leave the cost Hamiltonian constant. This is for example the case for trapped-ion platforms utilizing magnetic-gradient-induced-coupling (MAGIC), which we discussed in Chapter 1. In this case there is an always present all-to-all interaction between qubits, into which an optimization problem can be encoded.

Numerous extensions and variations of the basic quantum annealing protocol have been proposed to improve performance in practice. Nonlinear annealing schedules redistribute the evolution time so as to slow down near small spectral energy gaps and speed up in more favorable regions [280, 281]. Reverse annealing initializes the system in a classical candidate state and temporarily increases quantum fluctuations (strength of the drive Hamiltonian) to enable local exploration before re-annealing, effectively implementing a quantum local search [282–284]. Inhomogeneous driving, where the transverse-field strengths vary from qubit to qubit, has been shown to yield improvements of several orders of magnitude in the probability of finding the solution state for certain problem classes [285, 286] (see also discussion in Sec. 5.2). Counterdiabatic (or shortcut-to-adiabaticity) protocols add auxiliary driving terms designed to suppress diabatic transitions to excited states, thereby enabling high-fidelity ground-state preparation at shorter annealing times [287–289].

A further strategy is to augment the annealing Hamiltonian with a so-called catalyst Hamiltonian $\hat{H}_{\text{cat}}$,

$$\hat{H}(s) = s\hat{H}_{\text{cost}} + (1 - s) \hat{H}_{\text{drive}} + \lambda(s) \hat{H}_{\text{cat}}, \quad (3.10)$$

where $\lambda(s)$ is a time-dependent coefficient that vanishes at the beginning and end of the protocol, $\lambda(0) = \lambda(1) = 0$, so that the catalyst is only active during the intermediate stage of the annealing sweep. Due to this choice the ground state at the beginning and end of the annealing sweep remains unchanged. For certain problems the catalyst Hamiltonian can be chosen in a specific way in order to permit significantly faster annealing sweeps [195]. We will discuss the role of catalyst Hamiltonians and their impact on annealing performance in more detail in Sec. 5.3.

3.1.3 Adiabatic theorem

According to the adiabatic theorem, the system remains in its instantaneous ground state during a quantum annealing sweep (and the final state thus corresponds to the desired

ground state of $\hat{H}_{\text{cost}}$), provided the Hamiltonian changes sufficiently slowly such that the adiabatic condition holds [185–188]:

$$t_{\text{sweep}} \gg \hbar \max_{s \in [0,1]} \frac{|\langle n(s) | d\hat{H}(s)/ds | 0(s) \rangle|}{\Delta E_{n0}^2(s)}, \quad \forall n \geq 1. \quad (3.11)$$

Here, $\Delta E_{n0}(s)$ denotes the energy difference between the n -th instantaneous excited state $|n(s)\rangle$ and the instantaneous (non-degenerate) ground state $|0(s)\rangle$.

It should be noted that adiabaticity may not strictly be necessary for a successful annealing sweep. In fact, the search for optimized (non-adiabatic) annealing schedules and tighter bounds for the minimal quantum annealing time is an active field of research, see for example Refs. [290–292].

An order-of-magnitude estimate of the sweep duration required for a successful adiabatic evolution can be extracted directly from the right-hand side of Eq. (3.11). In this work, we employ the following simplified sufficient condition for adiabaticity:

$$t_{\text{sweep}} \gg \hbar \frac{V}{\Delta E^2} \equiv T. \quad (3.12)$$

Here, $\Delta E = \min_{s \in [0,1]} \Delta E_{10}(s)$ denotes the minimum energy gap between the instantaneous ground state and the first excited state along the annealing path, while $V = \max_{s \in [0,1]} |\langle 1(s) | d\hat{H}(s)/ds | 0(s) \rangle|$ captures the maximum magnitude of the driving-induced coupling between these two levels (if the first excited state is degenerate, the maximum is taken over the degenerate manifold). We will refer to the right-hand side of Eq. (3.12) as the *adiabaticity time* T .

The adiabaticity time provides a useful figure of merit for quantifying the computational hardness of a given annealing protocol: a larger T implies that a slower sweep, and hence a longer coherent evolution, is needed to maintain adiabaticity. We note, however, that evaluating T requires knowledge of the instantaneous energy spectrum along the entire annealing path, which in practice must be obtained numerically, e.g., through exact diagonalization. This limits the applicability of T as a diagnostic tool to moderate system sizes.

For hard optimization problems, the minimum gap is expected to close exponentially with the problem size N ,

$$\Delta E = \epsilon e^{-\alpha N}, \quad (3.13)$$

with a characteristic energy scale ϵ and an exponent $\alpha > 0$, a hallmark of a first-order quantum phase transition [293]. By contrast, the matrix element V in the numerator of Eq. (3.12) scales at most with the norm of the Hamiltonian, i.e., as $\mathcal{O}(N)$ for extensive systems. The adiabaticity time is therefore dominated by the exponential closing of the gap and grows as $T \propto e^{2\alpha N}$, up to polynomial corrections. Since this scaling is governed by ΔE , the minimum gap can be used as the primary figure of merit for the difficulty of solving a given optimization problem via quantum annealing.

3.1.4 Characteristic energy scale

So far, the annealing Hamiltonians in Eqs. (3.8)-(3.10) have been written without an explicit physical energy scale: the couplings J_{ij} and local fields h_i^z in $\hat{H}_{\text{cost}}$ (and also higher-order coupling in the case of PUBO) refer to the optimization problem alone, and the overall normalization is conventional. To make contact with a physical implementation, we factor

out a global energy scale J (with units of frequency; $\hbar = 1$) and write the physical annealing Hamiltonian as

$$\hat{H}_{\text{phys}}(t) = J \hat{H}(t), \quad (3.14)$$

where $\hat{H}(t)$ is dimensionless. The constant J is fixed by the hardware platform: it is the maximum two- or higher-order interaction strength that can be physically realized.

A natural normalization convention is to fix the largest dimensionless coupling to unity. For the quadratic cost Hamiltonian of Eq. (3.4) (QUBO), this corresponds to

$$\max_{ij} |J_{ij}| = 1, \quad (3.15)$$

so that J coincides exactly with the maximum two-body coupling rate of the device. By construction, all remaining dimensionless absolute couplings are then bounded by unity and lie automatically within the physical capabilities of the platform. The same applies to the longitudinal fields h_i^z , which on most hardware platforms, including MAGIC, can be implemented at rates well above the engineered spin-spin couplings.

The energy scale J also sets the time scale of the annealing dynamics. The physical sweep duration is

$$T_{\text{phys}} = t/J, \quad (3.16)$$

where t is the dimensionless sweep duration appearing in the annealing schedule. A larger value of J therefore translates directly into a proportionally shorter physical anneal, which makes the maximization of J within the hardware constraints a primary objective in the design of an analog quantum annealer.

For PUBO problems, the cost Hamiltonian generally contains k -body terms of different orders, which we discuss in detail in Chapter 5. On essentially all current quantum hardware, the rates associated with k -body interactions decrease with the interaction order k . On the MAGIC platform considered in Chapter 1, the higher-order spin-spin couplings arise as perturbative corrections to the leading two-body interaction and are therefore intrinsically weaker, as derived in Chapter 2. On digital platforms, multi-qubit operations require deeper gate compilations than two-qubit gates and consequently take longer to execute. The highest-order interaction present is therefore always the slowest one and sets the bottleneck for the achievable physical anneal time. The dimensionless cost Hamiltonian should accordingly be normalized to this term,

$$\max |J^{(k_{\text{max}})}| = 1, \quad (3.17)$$

so that J corresponds to the maximum physical rate of k_{max} -body interactions. Lower-order couplings, being implementable at comparatively higher rates, do not impose additional constraints on the schedule.

The driver strength h^x in $\hat{H}_{\text{drive}} = -h^x \sum_i \hat{\sigma}_i^x$ plays a fundamentally different role from J . Although h^x is also bounded from above by hardware limitations, the driving field is in practice implemented through a global laser or microwave pulse, whose Rabi frequency can be made significantly larger than any engineered spin-spin coupling. To good approximation, h^x can therefore be regarded as a free parameter of the annealing schedule rather than as a hardware-fixed energy scale. For the symmetric protocol of Eq. (3.8), a natural choice is $h^x = J$, that is, $h^x = 1$ in dimensionless units, which ensures that the driver and the cost Hamiltonian contribute with comparable weight along the sweep. We motivate this choice in the next section. For the modified protocol of Eq. (3.9), in which $\hat{H}_{\text{cost}}$ is present at all times, h^x instead enters as an additional optimization parameter whose optimal value depends on the problem instance and on the annealing schedule, as discussed in Chapter 4.




Figure 3.1: Influence of the driving strength on the minimum energy gap and the adiabaticity time for a linear annealing schedule $(1-s)h_x\hat{H}'_{\text{drive}} + sJ\hat{H}'_{\text{cost}}$. The cost Hamiltonian $\hat{H}_{\text{cost}} = J\hat{H}'_{\text{cost}}$ corresponds to the PUBO encoding of a random 3-SAT instance generated by uniquePT1 with $N = 6$ variables and has characteristic energy scale $J = \max_{ijk} |J_{ijk}^{(3)}|$. The driving Hamiltonian $\hat{H}_{\text{drive}} = h^x\hat{H}'_{\text{drive}} = -h^x \sum_i \hat{\sigma}_i^x$ represents a uniform transverse field of strength h^x . The minimum energy gap ΔE increases monotonically with the driving strength (red curve), while the adiabaticity time T , corresponding to the right-hand side of (3.12), exhibits a minimum at $h_x/J \approx 1.05$ (blue curve).

3.1.5 Optimal driving Hamiltonian strength

The strength h^x of the driving Hamiltonian H_{drive} is a free parameter that can be tuned to optimize the performance of the annealing schedule. In what follows, we discuss how the choice of the driving strength h^x influences the size of the minimum energy gap and the annealing performance for the symmetric annealing protocol defined in Eq. (3.8).

To this end, we introduce the dimensionless cost Hamiltonian $\hat{H}'_{\text{cost}} = \hat{H}_{\text{cost}}/J$ and the dimensionless driving Hamiltonian $\hat{H}'_{\text{drive}} = \hat{H}_{\text{drive}}/h^x$, expressed in units of the characteristic energy scale J and driving strength h^x , respectively. Consider now the dimensionless Hamiltonian $\hat{H}'(g) = \hat{H}'_{\text{cost}} + g\hat{H}'_{\text{drive}}$ and let us assume that the energy gap between the ground state and the first excited state of $\hat{H}'(g)$ becomes minimal at the (*a priori* unknown) critical value g_c . Writing the instantaneous Hamiltonian for the linear annealing schedule in Eq. (3.8) as $\hat{H}(s) = sJ\hat{H}'_{\text{cost}} + (1-s)h^x\hat{H}'_{\text{drive}}$, the relative strength g_c between $\hat{H}_{\text{cost}}$ and $\hat{H}_{\text{drive}}$, and thus the minimum energy gap, is reached at the time $s_c = (h^x/J)/(g_c + h^x/J)$ during the annealing sweep. The instantaneous Hamiltonian at this time reads $\hat{H}(s_c) = s_cJ\hat{H}'(g_c)$. Consequently, the minimum energy gap ΔE of the dimensionful Hamiltonian $\hat{H}$ relates to the dimensionless gap $\Delta E'_c$ of $\hat{H}'(g_c)$ as

$$\Delta E/J = s_c \Delta E'_c = \frac{h^x/J}{g_c + h^x/J} \Delta E'_c. \quad (3.18)$$

This relation shows that the minimum energy gap increases monotonically with the driving strength h_x and saturates at $J\Delta E'_c$ for large h_x/J , as illustrated in Fig. 3.1 for a single 3-SAT instance (see Sec. 3.2.5 for a discussion of the 3-SAT optimization problem).

Despite the monotonic behavior of the minimum energy gap, a larger driving strength does not necessarily result in a better annealing performance. This is because a large value of

h_x/J results in a steep ramp, making the schedule less adiabatic. Formally, this can be seen by examining the adiabatic condition in Eq. (3.12). As shown in Fig. 3.1 for a single 3-SAT instance, the adiabaticity time reaches a minimum where h_x and J become of comparable strength.

To understand this behavior, it is instructive to write the adiabatic condition in Eq. (3.12) as

$$t_{\text{sweep}} \gg T = \hbar \frac{V}{\Delta E^2} \approx \hbar \frac{J}{\Delta E^2} \left(V'_{\text{cost}} + \frac{h^x}{J} V'_{\text{drive}} \right). \quad (3.19)$$

Here, we have used the triangle inequality to separate the contributions of the cost and driving Hamiltonian to the quantity V into $V'_{\text{cost}} = \max_s |\langle 1(s) | \hat{H}'_{\text{cost}} | 0(s) \rangle|$ and $V'_{\text{drive}} = \max_s |\langle 1(s) | \hat{H}'_{\text{drive}} | 0(s) \rangle|$, assuming a linear annealing schedule. This approximation becomes exact in the limit of either small or large h^x/J . For small h_x/J , the contribution of V'_{drive} can be neglected and the adiabaticity time is dominated by the behavior of the minimum energy gap according to Eq. (3.18), causing the adiabaticity time to diverge for $h_x/J \rightarrow 0$. For large h_x/J , the adiabaticity time is dominated by the linearly increasing second term in Eq. (3.19). The optimal driving strength corresponds to a tradeoff between these two scenarios.

As this discussion shows, for a linear annealing schedule, the driving strength h^x should be chosen of comparable magnitude as the characteristic energy scale J of the cost Hamiltonian. The optimal value for h_x depends on (generally unknown) problem characteristics. For the specific instance used in Fig. 3.1, we find an optimal value of around $h_x/J \approx 1.05$. We checked that also for other problem instances the generic choice $h_x = J$ used throughout this work for the standard linear annealing schedule is not far from the optimal condition for adiabaticity.

However, we emphasize again that the discussion in this section only applies to the annealing schedule in Eq. 3.8. For the modified schedule in Eq. 3.9, the driving strength should generally be chosen larger than J , with the optimal value being highly problem dependent (see Chapter 4).

3.2 Benchmark models and combinatorial optimization problems

This section serves as a reference for the problems that will be revisited as benchmarks and test cases in the remainder of this thesis. The first two examples, the ferromagnetic p -spin model (Sec. 3.2.1) and a geometrically local Ising model of Ref. [294] (Sec. 3.2.2), are not classical optimization problems in the usual sense, but rather many-body Hamiltonians with specific properties, due to which they serve as valuable benchmarks for quantum annealing: their amenability to analytical or symmetry-reduced numerical treatment makes them ideal testbeds for studying spectral properties, finite-size scaling, and the role of non-stoquastic protocols. We then turn to two problems of direct industrial relevance, multiple object tracking (MOT) (Sec. 3.2.3) and the one-dimensional cutting stock problem (Sec. 3.2.4), for which we construct minimal QUBO instances that are tractable on near-term hardware. The section continues with 3-SAT (Sec. 3.2.5), the paradigmatic NP-complete problem of computer science, which we use as a benchmark for higher-order interactions in the cost Hamiltonian via its natural PUBO encoding. We close in Sec. 3.2.6 with a brief survey of further problems that admit a natural PUBO formulation, included for completeness.

3.2.1 p -spin model

A paradigmatic cost Hamiltonian frequently employed as a benchmark for quantum annealing is the ferromagnetic p -spin model ($p \geq 2$), defined by

$$\hat{H}_{\text{cost}} = -N \left(\frac{1}{N} \sum_{i=1}^N \hat{\sigma}_i^z \right)^p. \quad (3.20)$$

The ground state of this Hamiltonian is trivially identified: for odd p it is the fully polarized state with all spins pointing up, while for even p both fully polarized states with all spins up and all spins down are degenerate ground states. For $p = 2$, the model reduces to the Lipkin-Meshkov-Glick (LMG) model [295–297]. In the limit $p \rightarrow \infty$, the energy landscape becomes flat for all states except the ground state(s), and the model effectively recovers the Grover search problem [192].

Despite its trivial solution, the p -spin model exhibits rich behavior when subjected to a transverse-field driving Hamiltonian, making it a valuable testbed for quantum-optimization algorithms. Its fully connected, all-to-all interaction structure renders mean-field theory exact in the thermodynamic limit [192, 195, 298], allowing for analytical treatment via semiclassical methods. The standard quantum annealing protocol with a transverse-field driver encounters a first-order quantum phase transition for $p \geq 3$, accompanied by an exponentially small minimum energy gap [192]. This exponential gap closure implies that an adiabatic protocol would require a runtime growing exponentially with system size N .

Importantly, the p -spin model has served as the test problem for several strategies aimed at circumventing such first-order transitions in quantum annealing protocols. Non-stoquastic annealing Hamiltonians have been shown to soften the first-order transition to a second-order one for suitably designed protocols [195]. Since the gap at a second-order transition closes only polynomially with system size, this constitutes an exponential enhancement of annealing performance [196]. We will discuss non-stoquastic annealing protocols in more detail in Chapter 6.

The key feature enabling efficient numerical treatment of the p -spin model is its permutation symmetry: since the cost Hamiltonian (6.6) depends only on the collective magnetization $\sum_i \hat{\sigma}_i^z$, it commutes with all pairwise spin permutation operators. The same holds for the standard transverse-field driver $\hat{H}_{\text{drive}} = -\sum_i \hat{\sigma}_i^x$, and therefore for the total time-dependent annealing Hamiltonian. Consequently, the dynamics can be restricted to the fully symmetric subspace of the N -qubit Hilbert space.

This subspace is spanned by the Dicke states $|S, m\rangle$ with $S = N/2$, which are the simultaneous eigenstates of $\hat{\mathbf{S}}$ and $\hat{S}^z$, with the collective spin operators defined as

$$\hat{S}^\alpha = \frac{1}{2} \sum_{i=1}^N \hat{\sigma}_i^\alpha, \quad \alpha \in \{x, y, z\}, \quad (3.21)$$

satisfying $\hat{\mathbf{S}}^2 |S, m\rangle = S(S+1) |S, m\rangle$ and $\hat{S}^z |S, m\rangle = m |S, m\rangle$, where $m = -S, -S+1, \dots, S$. Each Dicke state $|N/2, m\rangle$ is the equal-weight symmetric superposition of all computational basis states with exactly $k = N/2 + m$ spins pointing up,

$$|N/2, m\rangle = \binom{N}{k}^{-1/2} \sum_{\substack{\text{permutations with} \\ k=N/2+m \text{ spins up}}} |\uparrow \cdots \uparrow \downarrow \cdots \downarrow\rangle. \quad (3.22)$$

Expressed in the Dicke basis, the cost Hamiltonian of the p -spin model takes the form

$$\hat{H}_{\text{cost}} = -N \left(\frac{2\hat{S}^z}{N} \right)^p = -\frac{2^p}{N^{p-1}} (\hat{S}^z)^p, \quad (3.23)$$

while the transverse-field driving Hamiltonian becomes ($h_x = 1$)

$$\hat{H}_{\text{drive}} = -2\hat{S}^x. \quad (3.24)$$

The action of the driver on the Dicke states can be calculated using the collective raising and lowering operators $\hat{S}^\pm = \hat{S}^x \pm i\hat{S}^y$:

$$\hat{S}^x |S, m\rangle = \frac{1}{2} \left(\sqrt{S(S+1) - m(m+1)} |S, m+1\rangle + \sqrt{S(S+1) - m(m-1)} |S, m-1\rangle \right). \quad (3.25)$$

The projection into the permutation-invariant subspace of the total Hilbert space then reduces the dimension from 2^N to $N+1$, enabling exact numerical diagonalization for systems of hundreds of qubits.

For even values of p , the Hamiltonian possesses an additional symmetry: the parity (or spin-flip) operator

$$\hat{\Pi} = \prod_{i=1}^N \hat{\sigma}_i^x, \quad (3.26)$$

which flips all spins simultaneously. Its action on the Dicke states is $\hat{\Pi} |S, m\rangle = |S, -m\rangle$, allowing the construction of parity eigenstates

$$|S, m\rangle_\pm = \frac{1}{\sqrt{2}} \left(|S, m\rangle \pm |S, -m\rangle \right), \quad (3.27)$$

for $m > 0$ and the $|S, 0\rangle_+ = |S, 0\rangle$ state belongs exclusively to the positive-parity sector. Working within a definite parity sector further reduces the Hilbert space dimension to $(N+1)/2$ for odd N , and $N/2 + 1$ for even N , the extra state in the latter case being the unpaired Dicke state $|S, 0\rangle$. Notably, the standard quantum annealing initial state $|+\rangle^{\otimes N}$, being the ground state of the transverse-field driving Hamiltonian, satisfies $\hat{\Pi} |+\rangle^{\otimes N} = +|+\rangle^{\otimes N}$ and thus belongs to the positive-parity sector (as well as the permutation-invariant subspace). Since $[\hat{\Pi}, H(t)] = 0$ throughout the annealing protocol (for even p), the dynamics remains confined to the positive-parity subspace, and it is this sector alone that needs to be diagonalized for even p .

3.2.2 Geometrically local Ising model

Here we discuss a specific geometrically local Ising model constructed by Albash in Ref. [294]. This model was designed such that a standard transverse-field annealing protocol develops a minimum energy gap that closes exponentially with system size, arising from a perturbative level crossing near the end of the anneal. Remarkably, augmenting the schedule with a suitably chosen non-stoquastic catalyst was reported to split this crossing into a sequence of avoided crossings, potentially yielding a polynomially closing minimum gap [294]. The model therefore serves as a natural benchmark for the role of non-stoquasticity in quantum annealing (see Sec. 5.3), a question we will discuss in detail in Chapter 6. Compared to the p -spin model of the previous section, the local Ising model is also closer to typical experimental settings, as it lacks permutation symmetry, involves only two-body interactions, and is defined on a graph with geometrically local connectivity.




Figure 3.2: Graph of the cost Hamiltonian H_{cost} of Eq. (3.28) for system sizes $N = 8, 10$, and 12 , following Ref. [294]. Each graph consists of two rings of $N/2$ qubits connected by two inter-ring bonds at diametrically opposite positions. Solid gray lines denote ferromagnetic couplings $J_{ij} = -1$, dashed black lines ferromagnetic couplings $J_{ij} = -1/2$, and dashed red lines antiferromagnetic couplings $J_{ij} = +5/6$ on the two inter-ring bonds. The site labels define the numbering convention used in our numerics.

Geometry and cost Hamiltonian. The model is defined on a graph consisting of two rings of $N/2$ qubits each (the model is only well-defined for even N), joined by two inter-ring bonds at diametrically opposite positions, as illustrated in Fig. 3.2 for $N = 8, 10$, and 12 . We label the sites $i = 0, \dots, N-1$, with the upper ring comprising sites $0, \dots, N/2-1$ and the lower ring comprising sites $N/2, \dots, N-1$, arranged such that site i and site $i + N/2$ sit at mirror-symmetric positions on the two rings. The cost Hamiltonian reads

$$H_{\text{cost}} = \sum_{\langle i,j \rangle} J_{ij} \hat{\sigma}_i^z \hat{\sigma}_j^z, \quad (3.28)$$

where the sum runs over nearest-neighbor pairs $\langle i,j \rangle$ on the graph. The couplings take three distinct values, indicated by different line styles in Fig. 3.2: ferromagnetic bonds with $J_{ij} = -1$ (solid gray), weakened ferromagnetic bonds with $J_{ij} = -1/2$ (dashed black), and antiferromagnetic bonds with $J_{ij} = +5/6$ on the two inter-ring links (dashed red). This coupling pattern was engineered in Ref. [294] so that two competing low-energy classical configurations become nearly degenerate, which is the mechanism responsible for the perturbative crossing mentioned above.

Discrete symmetries. Although H_{cost} is not permutation-symmetric, it is invariant under two commuting $\mathbb{Z}_2$ symmetries. The first is the global spin flip

$$\hat{\Pi} = \prod_{i=0}^{N-1} \hat{\sigma}_i^x, \quad (3.29)$$

which commutes with H_{cost} because each term in Eq. (3.28) contains an even number of Pauli Z operators. The second is the exchange of the upper and lower rings, implemented by the *ring-swap* operator $\hat{P}_{\text{swap}}$ that permutes sites i and $i + N/2$ for all $i = 0, \dots, N/2 - 1$. The graph and its coupling pattern are manifestly invariant under this exchange, so $[\hat{P}_{\text{swap}}, H_{\text{cost}}] = 0$.

It is convenient to write a computational basis state as $|a, b\rangle$ with $a, b \in \{0, 1\}^{N/2}$ denoting the bitstrings on the upper and lower ring, respectively. The two symmetry operators then act as

$$\hat{\Pi}|a, b\rangle = |\bar{a}, \bar{b}\rangle, \quad (3.30)$$

$$\hat{P}_{\text{swap}}|a, b\rangle = |b, a\rangle, \quad (3.31)$$

where $\bar{a}$ denotes the bitwise complement of a . Both operators square to the identity and commute with each other, generating the abelian group

$$G = \{1, \hat{\Pi}, \hat{P}_{\text{swap}}, \hat{\Pi}\hat{P}_{\text{swap}}\} \cong \mathbb{Z}_2 \times \mathbb{Z}_2. \quad (3.32)$$

The driver Hamiltonian used in the standard annealing schedule (Eq. (3.7)) is also invariant under G , so G is a symmetry of the full time-dependent Hamiltonian throughout the annealing protocol. In Chapter 6 we will extend the standard annealing protocol with a so-called catalyst term (see also Sec. 5.3) to study non-stoquastic annealing. We will see that this catalyst fulfills the same symmetries of the local Ising model, so that the discussion in this section remains applicable.

Reduction to the symmetric sector. Since G commutes with the full Hamiltonian at all times, the Hilbert space decomposes into invariant sectors labeled by the simultaneous eigenvalues $(\pm 1, \pm 1)$ of $(\hat{\Pi}, \hat{P}_{\text{swap}})$. The annealing protocol starts from the initial state $|+\rangle^{\otimes N}$, which is the uniform superposition over all computational basis states and therefore a $+1$ eigenstate of both symmetry operators. Consequently, the dynamics remains confined to the $(+1, +1)$ sector, which we refer to as the *symmetric sector*, and it is sufficient to diagonalize the Hamiltonian within this subspace.

An orthonormal basis of the symmetric sector is obtained by partitioning the 2^N computational basis states into orbits under the action of G and assigning, to each orbit $\mathcal{O}$, the normalized equal superposition

$$|\mathcal{O}\rangle = \frac{1}{\sqrt{|\mathcal{O}|}} \sum_{|s\rangle \in \mathcal{O}} |s\rangle. \quad (3.33)$$

Each $|\mathcal{O}\rangle$ is manifestly invariant under every element of G , orbits with disjoint supports yield orthogonal states, and any vector in the symmetric sector can be expanded as a linear combination of such orbit states. The dimension d of the symmetric sector is therefore equal to the number of orbits, which, by Burnside's lemma, equals the average number of computational basis states fixed by the elements of G .

A direct enumeration of fixed points gives the following results. The identity fixes all 2^N basis states. The spin flip $\hat{\Pi}$ fixes no state, since no bitstring is equal to its own complement. The ring-swap $\hat{P}_{\text{swap}}$ fixes the $2^{N/2}$ states satisfying $a = b$, and the combined operation $\hat{\Pi}\hat{P}_{\text{swap}}$ fixes the $2^{N/2}$ states satisfying $a = \bar{b}$. Burnside's lemma then yields

$$d = \frac{1}{|G|} \sum_{g \in G} |\text{Fix}(g)| = \frac{1}{4} (2^N + 0 + 2^{N/2} + 2^{N/2}) = 2^{N-2} + 2^{N/2-1}, \quad (3.34)$$

which approaches $2^N/4$ at large N . Although the reduction is only by a constant factor, so that the symmetric-sector dimension remains exponential in N , it extends the range of system sizes accessible to exact diagonalization by around two qubits, which is still useful for studying finite-size scaling effects.

3.2.3 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 [299]. 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) [300], 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 [301].

In the context of quantum annealing, focusing on the challenge of associating given detections to coherent tracks, a QUBO formulation offers a suitable framework. To construct the QUBO instance, we follow the approach discussed in Ref. [302], which scales linearly with the number of detections, tracks, and frames and could, in principle, scale to industry-relevant problem sizes. As this will serve as an initial test problem for practical quantum annealing protocols in the presence of noise in Chapter 4, we introduce simplifying assumptions to reduce the problem size. A related work, which focuses on fast, online algorithms, is presented in Ref. [303].

General QUBO formulation



Figure 3.3: Two frames of a video from Refs. [304, 305] 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.

The QUBO formulation of the MOT 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

$$x^{(f)} \in \{0, 1\}^{D_f T}, \quad (3.35)$$

which encodes detections as

$$x_{k.T+l}^{(f)} = 1 \iff \text{Detection } k \text{ in frame } f \text{ is in track } l. \quad (3.36)$$

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:

$$\sum_{l=0}^{T-1} x_{k.T+l}^{(f)} = 1 \quad \forall k \in \{0, \dots, D-1\}, \quad (3.37)$$

$$\sum_{k=0}^{D-1} x_{k.T+l}^{(f)} = 1 \quad \forall l \in \{0, \dots, T-1\}. \quad (3.38)$$

Figure 3.3 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 bounding boxes B_1 and B_2 :

$$G(B_1, B_2) := \frac{|B_1 \cap B_2|}{|B_1 \cup B_2|}. \quad (3.39)$$

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 pre-trained neural network [306]:

$$C(f_{k_1}, f_{k_2}) := \frac{f_{k_1}^T f_{k_2}}{\|f_{k_1}\| \cdot \|f_{k_2}\|}. \quad (3.40)$$

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. [307].

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

$$W_{m,l} = \frac{1}{2}C(f_{k_1}, f_{k_2}) + \frac{1}{2}G(B_{k_1}, B_{k_2}). \quad (3.41)$$

To combine the objective and the constraints (Eqs. (3.37) and (3.38)) 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

$$\begin{aligned} & \arg \min_{\{x \in \{0,1\}^n \mid Ax=b\}} x^T W x \\ &= \arg \min_{\{x \in \{0,1\}^n\}} x^T W x + \lambda \|Ax - b\|^2. \end{aligned} \quad (3.42)$$

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.

Minimal MOT instance with four qubits

We construct the smallest non-trivial 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 8-dimensional ($\mathbf{x} = (\mathbf{x}^{(1)}, \mathbf{x}^{(2)})^T$, see Eq. (3.35)). 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. (3.37) and (3.38), 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 $\mathbf{x}^{(2)}$. Formulating the constraints as a linear system of equations $A\mathbf{x}^{(2)} = \mathbf{b}$, we get

$$\begin{pmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \end{pmatrix} \mathbf{x}^{(2)} = \begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}. \quad (3.43)$$

Since we fixed the assignments of detections to tracks on the first frame, the objective part W of the cost function in Eq. (3.42) 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. 3.3, the objective is given by $W \approx (-2.18, -0.95, -0.58, -2.26)$.

To write Eq. (3.42) as a cost Hamiltonian for quantum annealing, we translate the Boolean variables to quantum spins via the transformation $x_i \rightarrow (1 + \sigma_i^z)/2$ (see Sec. 3.1.1). The penalty factor λ in Eq. (3.42) has a large effect on the minimal energy gap in the corresponding quantum annealing protocol and thus has to be chosen carefully [203, 308]. 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. [309]). Finally, when utilizing the minimal MOT instance for our numerical simulations in Chapter 4, 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. 3.1.2.

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

$$H_{\text{cost}}/J = \begin{pmatrix} 0 & -0.87 & -0.38 & -0.23 & -0.91 \\ 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}, \quad (3.44)$$

which we write in units of the relevant characteristic energy scale, i.e., the maximal realizable two-body coupling strength $J := \max_{ij} |J_{ij}|$ (see discussion in Sec. 3.1.4). This energy scale




Figure 3.4: Final fidelities achieved for different total linear annealing sweep durations for the MOT QUBO problem with 9 qubits. We consider the modified annealing protocol defined in Eq. 3.9. 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. (3.42)) as well as the strength of the driver Hamiltonian h_x/J (Eq. (3.7)). The vertical red dashed line denotes the sweep time $2.6J^{-1}$ used in our numerical simulations in Chapter 4. In an intermediate regime of medium-slow sweeps, it can be beneficial to carefully tune these hyperparameters, either through a variation of test runs or heuristic methods for finding optimal penalty parameters.

depends on the considered hardware platform. We discussed details for trapped-ion hardware using magnetic gradient-induced couplings in Chapter 1.

MOT instance with nine qubits

We now construct the next largest MOT QUBO instance, which requires 9 qubits in our encoding (see Sec. 3.1.2). The 9-qubit instance is generated by adding a third frame to the video sequence used for the minimal problem (Fig. 3.3). Following the same pattern as before, we consider the first, third, and fifth frame of the video, which yields a 9-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. (3.41)). For this particular instance we observe slightly increased annealing performance with this rescaling and the remaining chosen parameters (see below). The resulting 9×9 QUBO matrix, where entries equal to one correspond to constraint terms and the remaining entries encode the objective function, is

$$H_{\text{cost}}/J = \begin{pmatrix} 0 & -0.96 & -0.55 & -0.43 & -0.99 & -0.84 & -0.89 & -0.38 & -0.97 \\ 0 & 0 & 1 & 1 & 0 & -0.19 & 0 & -0.04 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & -0.19 & 0 & -0.04 \\ 0 & 0 & 0 & 0 & 1 & -0.05 & 0 & -0.18 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -0.05 & 0 & -0.18 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}. \quad (3.45)$$

The optimal penalty factor λ for the QUBO problem (see Eq. (3.42)) varies between problem instances and must be determined individually. For this 9-qubit problem, we select $\lambda = 3$, which provides good annealing performance at the reference time of $2.6J^{-1}$ used throughout Chapter 4. Similarly, the optimal transverse driving strength h_x (Eq. (3.7)) differs from the minimal problem: while the 5-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 3.4 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 red vertical dashed 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 [309].

For the numerical simulations in Chapter 4, we employ $h_x/J = 2$ and $\lambda = 3$ without further optimization, as these parameters yield sufficiently good results for the primary objective of demonstrating the efficacy of dynamical decoupling for quantum annealing with tractable instances of industrially-relevant optimization problems.

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. (3.35). 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 further away than the maximal frame gap, a new person enters the frame, then one can use the same ID [302]. The two persons are then separated in a post-processing phase. Some benchmark results showing how classical solvers perform on the more challenging MOT20 dataset are given in Ref. [310].

3.2.4 Cutting stock problem

In this section we give a brief overview over the industrially relevant cutting stock problem [311], which will be used in addition to MOT as a benchmark problem in Chapter 4. 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

$$\arg \max_{x \in \mathbb{Z}^N} \sum_{i=1}^m \sum_{k=1}^Q l_i x_{i,k}, \quad (3.46)$$

with the constraints

$$\sum_{i=1}^m l_i x_{i,k} \leq L \quad \forall k \in \{1, \dots, Q\}, \quad (3.47)$$

$$\sum_{k=1}^Q x_{i,k} \leq d_i \quad \forall i \in \{1, \dots, m\}. \quad (3.48)$$

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 [312–314].

An important step for the reformulation of Eq. (3.46) into a QUBO problem is to turn all inequality constraints into equality constraints, which can be achieved by using slack variables s :

$$Ax - b \leq 0 \Leftrightarrow Ax - b = s. \quad (3.49)$$

The slack variables are non-positive and are restricted to values

$$s_i \in \mathcal{S}_i := \{(Ax - b)_i | x \in \mathcal{F}\}, \quad (3.50)$$

where $\mathcal{F}$ is the feasible set for x . One usually has to implement integer s_i variables via a binary representation

$$s_i = - \sum_{k=0}^{B_{\max}} 2^k s_i^{(k)}, \quad (3.51)$$

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. (3.42). Lower bounds for the penalty parameters of the multi-knapsack problem have already been computed, which can be directly applied to our setting [315]. Alternatives to implement inequalities without additional slack variables are, e.g., the unbalanced penalization formulations [316] or by the use of an ancillary qudit in discretized quantum optimization schemes [317].

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. (3.47), we would need $1 + \lfloor \log_2(L) \rfloor$ slack variables for every bar stock and because of the inequality given by Eq. (3.48), $1 + \lfloor \log_2(d_i) \rfloor$ slack variables for every demanded piece i . However, the latter inequalities can also be incorporated without slack variables as we will show now.

For problem instances with multiple larger bars, the inequalities in Eq. (3.48) become nontrivial. In the case with two bars, they read

$$x_{1,1} + x_{1,2} \leq 1, \quad (3.52)$$

$$x_{2,1} + x_{2,2} \leq 1. \quad (3.53)$$

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. [318]. The same procedure can also be generalized for larger numbers of bars by adding

$$\sum_{i=1}^m \left(\left(\sum_{k=1}^Q x_{i,k} \right) \left(\sum_{k=1}^Q x_{i,k} \right) - \sum_{k=1}^Q x_{i,k}^2 \right). \quad (3.54)$$

Minimal cutting stock problem instances

For benchmarking purposes, it is useful to construct 1D cutting stock problems that have a low qubit requirement. We explicitly construct two trivial cutting stock problems requiring four and five binary variables, respectively, which we use in Chapter 4 to test a dynamical decoupling protocol. The problems are illustrated in Fig. 3.5. 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. (3.48) are trivially fulfilled, since there is only one larger bar stock. There is only one length constraint (Eq. (3.47)). It reads

$$1x_1 + 1x_2 \leq 3. \quad (3.55)$$

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

$$(1 \quad 1 \quad 1 \quad 2) \begin{pmatrix} x_1 \\ x_2 \\ s_1 \\ s_2 \end{pmatrix} = (3). \quad (3.56)$$

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. (3.56) is denoted as $Ax = b$. This calculation results in the matrix

$$Q_{\text{const.}} = \lambda \begin{pmatrix} -5 & 1 & 1 & 2 \\ 1 & -5 & 1 & 2 \\ 1 & 1 & -5 & 2 \\ 2 & 2 & 2 & -8 \end{pmatrix}. \quad (3.57)$$

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




Figure 3.5: Illustration of the small cutting stock toy problems. a) and b) show problem instances which require four and five qubits, respectively. When encoding the problem into two-body interactions only, an additional ancilla qubit is required (Eq. (3.6)). In both cases, the trivial optimal solution is that every required piece will be cut out from the larger bar.

$$H_{\text{cost}}/J = \begin{pmatrix} 0 & -1 & -1 & -0.5 & -1 \\ 0 & 0 & 0.5 & 0.5 & 1 \\ 0 & 0 & 0 & 0.5 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}. \quad (3.58)$$

The slightly larger problem which is constructed with 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 $1 + \lceil \log_2(4) \rceil = 3$ slack variables, resulting in the equality:

$$(2 \quad 2 \quad 1 \quad 2 \quad 4) \begin{pmatrix} x_1 \\ x_2 \\ s_1 \\ s_2 \\ s_3 \end{pmatrix} = (4). \quad (3.59)$$

In the same way as in Eq. (3.57), this equality yields the contribution

$$Q_{\text{const.}} = \lambda \begin{pmatrix} -12 & 4 & 2 & 4 & 8 \\ 4 & -12 & 2 & 4 & 8 \\ 2 & 2 & -7 & 2 & 4 \\ 4 & 4 & 2 & -12 & 8 \\ 8 & 8 & 4 & 8 & -16 \end{pmatrix} \quad (3.60)$$

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

$$H_{\text{cost}}/J = \begin{pmatrix} 0 & 0.33 & 0.33 & 0.25 & 0.5 & 1 \\ 0 & 0 & 0.33 & 1.67 & 0.33 & 0.67 \\ 0 & 0 & 0 & 1.67 & 0.33 & 0.67 \\ 0 & 0 & 0 & 0 & 1.67 & 0.33 \\ 0 & 0 & 0 & 0 & 0 & 0.67 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}. \quad (3.61)$$

3.2.5 3-SAT

In this section we give the background on the paradigmatic 3-SAT optimization problem, which will be utilized as a benchmark problem in Chapter 5, when discussing the influence of PUBO cost Hamiltonians for quantum annealing. Given a Boolean formula, the satisfiability problem (SAT) poses the decision problem whether there is an assignment of truth values such that the formula evaluates to *true*. In the case of k -SAT, the formula takes the form of a conjunction of M logical clauses C_m ,

$$C_1 \wedge C_2 \wedge \dots \wedge C_M, \quad (3.62)$$

with each clause being composed of a disjunction of k literals chosen from a set of N Boolean variables x_i and their negations $\neg x_i$. An example clause for $k = 3$ is

$$C_m(x_i, x_j, x_l) = x_i \vee \neg x_j \vee \neg x_l. \quad (3.63)$$

Every SAT problem can be rewritten as a k -SAT problem using the rules of Boolean algebra. It is well known that k -SAT for $k \leq 2$ belongs to the complexity class P, i.e., it is solvable in polynomial time on a classical computer [319]. In contrast, 3-SAT was the first decision problem proven to be NP-complete (Cook–Levin theorem [320, 321]). If $P \neq NP$, no classical algorithm can solve 3-SAT in polynomial time. Additionally, it is widely believed that NP-complete problems are not contained in BQP [322], and could therefore not be solved in polynomial time on a quantum computer either. Nevertheless, even if an exponential quantum advantage may be out of reach, one can still hope for a practical speedup, e.g. with quantum annealing [77, 183, 323–325].

A phase transition for 3-SAT problems is known to occur asymptotically at a critical ratio of clauses to variables of $(M/N)_{\text{crit}} \approx 4.24$ [326]. Around this critical point, one is likely to encounter hard-to-solve problems with unique solutions. Below the phase transition point, most problem instances are easy to solve with many solutions, while above the transition most instances become unsatisfiable. Note that although the probability of finding instances with unique solutions below the phase transition point decreases exponentially, it is well-known that those rare instances (with fewer clauses and unique solutions) tend to be the hardest ones to solve [327].

Generating 3-SAT instances

To generate specific instances of the 3-SAT problem, used in Chapter 5, we adapt the three different generators **toughSAT**, **uniquePT1**, and **uniquePT4** described in Ref. [328] (see also Sec. 5.1.1).

The **toughSAT** generator creates an arbitrary number M of 3-SAT clauses by randomly picking three out of N variables and then randomly applying negations to them. For our numerical simulations in Chapter 5, we set the ratio of clauses to variables close to the critical ratio $(M/N)_{\text{crit}} \approx 4.24$, as we expect to find hard-to-solve instances with unique ground states in the vicinity of this point. To avoid ground state degeneracies, we postselect the randomly generated instances for those with unique solutions.

In contrast, the generators **uniquePT1** and **uniquePT4** always create instances with known and unique solutions (although not necessarily hard to solve in the latter case). Both generators work by picking a solution randomly and then constructing appropriate 3-SAT clauses to make said solution unique. While **uniquePT4** creates $M = 4N$ clauses to achieve that, **uniquePT1** only needs $M = N + 6$ clauses. However, we find that **uniquePT4** constructs clauses in such a way that all third-order and second-order terms in the classical cost function, corresponding to three-body and two-body interactions in the PUBO cost Hamiltonian, cancel each other. What remains is a simple cost Hamiltonian with only local terms of the form $\hat{H}_{\text{cost}} = -\sum_i h_i^z \hat{s}_i^z$, whose solution can easily be read off. Thus, **uniquePT4** creates trivial 3-SAT problems, where the time required to solve them scales linearly in the problem size.

For our numerical analysis in Chapter 5, we are primarily interested in hard-to-solve problem instances generated by **toughSAT** and **uniquePT1**, which exhibit an exponential closing of the minimum energy gap with increasing problem size during an annealing protocol.

3-SAT as a PUBO problem

3-SAT can be reformulated in various ways as a PUBO or QUBO problem in order to make it suitable for quantum annealing. A straightforward encoding consists in a PUBO problem of degree three: each clause is cast into the form of an energy term penalizing any choice of truth values that does not satisfy the clause. For example, the clause $C_m = x_i \vee \neg x_j \vee \neg x_\ell$ from above can be translated to $(1 - x_i)x_j x_\ell$, which only evaluates to zero if at least one variable satisfies the corresponding clause. Adding up these terms for each clause and translating the binary variables to quantum spin variables via $x_i = 1/2 - \hat{s}_i^z$, leads to a general three-body spin Hamiltonian of PUBO form. We investigate such higher-order cost Hamiltonians and their effect on annealing performance in Chapter 5. Note that, in this specific PUBO formulation, the three-body and two-body interaction strengths $J^{(3)}$ and $J^{(2)}$ as well as the local biases h^z are all of the same order of magnitude or zero. The ground state of the cost Hamiltonian then encodes the solution of the original 3-SAT problem (more precisely, it solves MAX-SAT, i.e., finding the solution that satisfies most clauses). This encoding requires N spins (qubits) for the N Boolean variables in the original problem.

3-SAT as QUBO via slack variables

Currently available quantum annealing hardware (e.g., the D-Wave machine [329]) natively only supports two-body interactions. Thus, one needs to find an equivalent encoding as a QUBO problem to solve 3-SAT on such hardware. A resource-efficient way to achieve this is to start with the PUBO formulation of 3-SAT given by a general cost Hamiltonian with

three-body interactions, two-body interactions and local fields, and to introduce slack variables that reduce all three-body interactions to quadratic and local terms (we will explicitly demonstrate this process with a toy model in Sec. 5.1.1). Since there are at most M (number of clauses) unique three-body terms in the Hamiltonian, at most M additional qubits are required by this QUBO reduction. For hard problem instances around the phase transition this amounts to around $4.24 N$ additional qubits.

Finding an optimal QUBO reduction for a given PUBO problem is in general NP-hard. However, for small system sizes, we can employ the following brute-force algorithm to find the most efficient reduction: replace the pair of variables that appears most frequently in the third-order terms of the PUBO cost function by a slack variable, then continue with the second-most-frequent pair of variables and so on, until all third-order terms have been replaced by quadratic ones with associated slack variables and constraints [330]. For each newly introduced ancilla variable $y = x_i x_j$, a penalty term $\propto (3y + x_i x_j - 2x_i y - 2x_j y)$ needs to be added to the cost Hamiltonian and scaled by a suitable factor, thus ensuring that violating the penalty term in favor of the three-body interactions is never energetically favorable. This factor depends on the problem size and structure: in our numerics, we make the simple choice that for every three-body interaction term $a_{ijk} x_i x_j x_k$ that is to be reduced by the ancilla variable $x_{\text{ancilla}} = x_i x_j$, we add $|a_{ijk}| + 1$ to the strength of the corresponding penalty term. In Sec. 5.1.2 of Chapter 5, we discuss how varying the strength of the penalty terms relative to the original cost Hamiltonian affects the overall annealing performance.

The QUBO formulation of **toughSAT** instances usually requires more slack variables than **uniquePT1** instances of the same size. This is due to the fact that the number of ancilla variables is bounded from above by the number of clauses in the problem instance. For both generators, the increased resource requirements of the QUBO reductions generally lead to a smaller minimum gap and thus reduced annealing performance.

3-SAT as QUBO via the maximum independent set problem

A well-known example of a QUBO formulation of 3-SAT uses the equivalent maximal independent set problem (MIS) [331]: Draw a graph with three vertices for each clause (corresponding to the three involved literals). All vertices belonging to the same clause as well as pairs of vertices corresponding to mutually exclusive literals (i.e., x_i and $\neg x_i$) are connected by edges, the set of which is denoted with E . The MIS problem then asks whether it is possible to color M vertices without two colored vertices being connected by an edge. If the answer is yes, the original 3-SAT problem is satisfiable. This MIS problem can be easily translated to a QUBO problem requiring $3M$ qubits (or about $\text{round}(3 \times 4.24 N) = \text{round}(12.72 N)$ for hard problems):

$$f_{\text{MIS}}(\mathbf{x}) = \sum_{ij \in E} x_i x_j. \quad (3.64)$$

Although generically more resources are required than for the reduction with slack variables, a potential advantage of the MIS formulation, especially regarding its use on physical quantum devices, lies in the fact that all interactions have the same energy scale, whereas slack variables for a direct QUBO reduction are introduced with constraints that are imposed through high-energy penalties.

3-SAT as QUBO via linear inequality constraints

Finally, we mention also a linear formulation of 3-SAT as a QUBO problem [332], where one introduces one ancillary Boolean variable z_m for each clause C_m that indicates whether

said clause is fulfilled or not. One then asks to maximize the number of fulfilled clauses

$$\max \sum_m z_m \quad (3.65)$$

under the inequality constraints

$$\sum_{x_i \in \text{TRUE}_m} x_i + \sum_{x_i \in \text{FALSE}_m} (1 - x_i) \geq z_m, \quad (3.66)$$

where TRUE_m (FALSE_m) is the set of literals contained in C_m satisfying the clause if chosen as *true* (*false*). One can then introduce two slack variables per inequality constraint to express them as equalities and write the problem in QUBO form.

Overall, this approach needs $3M$ ancilla qubits in addition to the N qubits for the literals, as well as an energy penalty to enforce the constraints. This scaling has to be compared to the one using the mapping to the maximal independent set problem ($3M$ qubits, no different energy scales), the one decomposing the PUBO problem using slack variables (N qubits plus at most M ancillas, plus energy penalties), and the direct PUBO implementation (N qubits, interactions of degree three, no different energy scales).

3.2.6 Selected problems with natural PUBO formulations

There exists a plethora of other optimization problems that are naturally posed as PUBO. In this section, we mention a few illustrative examples.

NAE-k-SAT

Apart from 3-SAT, which we have discussed above, there are many other variations and constraints on the general satisfiability problem. One of those is the NAE- k -SAT (Not-All-Equal- k -SAT) problem [333], which can also be naturally formulated as a PUBO. Like k -SAT, NAE- k -SAT consists of a conjunction of clauses, where each clause contains k Boolean variables and is only *true*, if at least one variable is *true* and at least one variable is *false*. Only for $k \geq 4$ is the problem naturally formulated as a PUBO with higher-order interactions [333]. For example, for $k = 4$, a clause $C(x_i, x_j, x_k, x_l)$ can be encoded into one four-body interaction $\sim \hat{s}_i^z \hat{s}_j^z \hat{s}_k^z \hat{s}_l^z$ and six two-body interaction penalty terms of the cost Hamiltonian.

As for k -SAT, this PUBO encoding requires one qubit for each Boolean variable. The equivalent reduction to a QUBO, on the other hand, requires (for $k = 4$) in the worst case two additional ancilla qubits for each clause in the problem.

Hypergraph coloring

Consider an undirected graph $G = (V, E)$, consisting of a finite set of vertices V and a set of edges E , the latter being a set of subsets of V . For a regular undirected graph, all edges are unordered pairs of vertices. For a hypergraph on the other hand, each edge can contain any number of vertices.

Given G and a number of k colors, the hypergraph vertex coloring problem asks whether there is an assignment of exactly one colour to each vertex such that no edge contains vertices all of the same colour [334]. A QUBO formulation of the regular graph coloring is

well-known [189], which can be easily generalized to hypergraphs: let $x_n^i = 1$ if vertex n has color i , else 0. The cost function of the problem can then be written as

$$f(\mathbf{x}) = \sum_{n \in V} \left(1 - \sum_{i=1}^k x_n^i\right)^2 + \sum_{e \in E} \sum_{i=1}^k \prod_{n \in e} x_n^i, \quad (3.67)$$

where the first term ensures that each vertex has exactly one color assigned to it and the second term penalizes each edge whose vertices are all of the same color. Thus, any solution with zero energy is a proper vertex coloring of the hypergraph.

To implement this PUBO problem encoding on a quantum computer, $k|V|$ qubits are required (with $|V|$ being the number of vertices in the graph). To reduce Eq. (3.67) to a QUBO problem, one needs to introduce ancilla variables for every edge in the graph containing more than two vertices (the exact amount needed will depend on the degree of the edges and their overlap with each other).

Other generalizations of graph problems to hypergraphs similarly yield PUBO problems with higher-order interaction terms, e.g., the vertex cover problem or the graph partitioning problem [335].

Traveling salesperson problem with time windows

A variation on the famous traveling salesperson problem (TSP) is the traveling salesperson problem with time windows (TSPTW). Here, like in the TSP, we ask for a Hamiltonian path with minimum weight on a graph with weights assigned to all edges. In addition, each vertex has to be visited in a specific time window. Both QUBO and PUBO formulations of this problem are known. In Ref. [336], the authors propose two different QUBO and one PUBO formulation of degree four for the problem. In the PUBO case, the number of required qubits scales as $\mathcal{O}(n^2 + n\delta)$ in the number of vertices n (here, δ is a number scaling at most linearly in n). In contrast, the QUBO formulations scale as $\mathcal{O}(n^3 + n\delta)$ and $\mathcal{O}(n^2 + n^2\delta)$, respectively, for the edge-based and integer linear programming (ILP) formulations.

3.3 Related quantum algorithms for solving optimization problems

Quantum annealing was originally conceptualized for analog platforms, where the relative strength of the cost and driving Hamiltonians can be tuned continuously. With the increasing availability of digital quantum computers, gate-based variants have become popular alternatives. Digital implementations are particularly attractive for PUBO problems, as they allow higher-order interactions to be synthesized from universal sets of single- and two-qubit gates, whereas engineering native multi-qubit couplings in analog devices remains challenging (see discussions in Secs. 2.3 and 5.1.5). In this section, we introduce two closely related gate-based approaches to quantum annealing: Trotterized quantum annealing, which discretizes the continuous-time evolution into a sequence of quantum gates, and the Quantum Approximate Optimization Algorithm (QAOA) [63], a variational generalization that replaces the fixed annealing schedule with optimizable parameters.

3.3.1 Trotterized quantum annealing

Digital quantum annealing (or Trotterized quantum annealing) is accomplished via a Trotterization of the time evolution under the Hamiltonian in Eq. 3.8. That is, for a grid of times

$0 = t_0 < t_1 < \dots < t_M = t_{\text{sweep}}$ with uniform step size $\Delta t = t_{\text{sweep}}/M$, the time evolution operator can be approximated by decomposing it into a product of simpler exponentials.

The simplest approach is the first-order Lie–Trotter decomposition [337, 338], which yields

$$\hat{\mathcal{T}} e^{-it_{\text{sweep}} \int_0^1 ds' \hat{H}(s')/\hbar} = \prod_{m=1}^M e^{-i\Delta t(1-s_m)\hat{H}_{\text{drive}}/\hbar} e^{-i\Delta t s_m \hat{H}_{\text{cost}}/\hbar} + \mathcal{O}(\Delta t), \quad (3.68)$$

where $\hat{\mathcal{T}}$ is the time-ordering operator and $s_m = t_m/t_{\text{sweep}}$. The local error for each Trotter step is $\mathcal{O}(\Delta t^2)$, resulting in an accumulated global error that scales as $\mathcal{O}(\Delta t)$, which can be improved through higher-order splitting methods [339].

To that end, a more accurate approximation is provided by the second-order symmetric decomposition, also known as the midpoint rule [340]. This formula evaluates the time-dependent coefficients at the midpoint $s_{m-1/2} = (m-1/2)/M$ of each time interval and symmetrically splits the driving Hamiltonian:

$$\begin{aligned} \hat{\mathcal{T}} e^{-it_{\text{sweep}} \int_0^1 ds' \hat{H}(s')/\hbar} &= \prod_{m=1}^M e^{-i\Delta t(1-s_{m-1/2})\hat{H}_{\text{drive}}/(2\hbar)} e^{-i\Delta t s_{m-1/2} \hat{H}_{\text{cost}}/\hbar} \\ &\times e^{-i\Delta t(1-s_{m-1/2})\hat{H}_{\text{drive}}/(2\hbar)} + \mathcal{O}(\Delta t^2). \end{aligned} \quad (3.69)$$

This symmetric splitting achieves a local error of $\mathcal{O}(\Delta t^3)$ per step, leading to a global error of $\mathcal{O}(\Delta t^2)$. The improved error scaling implies that fewer Trotter steps are required to reach a target accuracy, which can substantially reduce the circuit depth in practice.

In practical implementations on near-term quantum hardware, single-qubit gates are typically much faster and exhibit significantly higher fidelities than two-qubit entangling gates. Since the driving Hamiltonian $\hat{H}_{\text{drive}}$ consists solely of single-qubit terms, splitting it symmetrically as in the second-order formula introduces only additional single-qubit rotations, leaving the number of two-qubit gates per Trotter step unchanged. Consequently, the second-order decomposition does not substantially increase the effective circuit complexity or depth in terms of the error-prone two-qubit operations. Given this negligible overhead and the quadratically improved error scaling, there is in practice little reason to employ the first-order formula.

Since all terms within the cost Hamiltonian commute with each other, the corresponding matrix exponential can be exactly decomposed into a product of single-qubit Z rotations and two-qubit ZZ rotations. Similarly, the matrix exponential of the driving Hamiltonian decomposes into independent single-qubit X rotations. Each two-qubit rotation may then be further compiled into the native gate set of the respective architecture.

Compared to analog quantum annealing, digital Trotterization incurs additional errors due to the time discretization. Furthermore, since each Trotter step must be decomposed into native gates, there is generally a polynomial overhead in the number of gates. For sufficiently small Trotter steps, however, the limiting factors to the algorithm’s performance are the same as those of the continuous protocol. In this regime, the adiabaticity time estimated from the minimum energy gap of the analog quantum annealing protocol serves as a proxy for the number of discrete time steps, and thus the circuit depth, required to reach the ground state of the cost Hamiltonian.

3.3.2 Quantum approximate optimization algorithm (QAOA)

The Quantum Approximate Optimization Algorithm (QAOA), introduced in Ref. [63], provides an alternative gate-based approach to combinatorial optimization that is closely related to quantum annealing. The hybrid classical-quantum algorithm produces a variational quantum state through alternating applications of two unitary operators: one encoding the cost function and one acting as a drive Hamiltonian (also called 'mixer') that causes transitions between computational basis states.

Using the notation for drive and cost Hamiltonian as established in the previous sections, the QAOA prepares the parameterized state

$$|\gamma, \beta\rangle = \prod_{m=1}^p e^{-i\beta_m \hat{H}_{\text{drive}}/\hbar} e^{-i\gamma_m \hat{H}_{\text{cost}}/\hbar} |+\rangle^{\otimes n}, \quad (3.70)$$

where p denotes the circuit depth, $\gamma = (\gamma_1, \dots, \gamma_p)$ and $\beta = (\beta_1, \dots, \beta_p)$ are classical variational parameters, and $|+\rangle^{\otimes n}$ is the uniform superposition over all computational basis states. The expectation value $\langle \gamma, \beta | \hat{H}_{\text{cost}} | \gamma, \beta \rangle$ is then optimized classically over the $2p$ parameters. Importantly, the quality of the approximation improves monotonically with increasing depth, and in the limit $p \rightarrow \infty$, the algorithm can produce exact solutions [63].

The structural similarity between QAOA and Trotterized quantum annealing is evident: both algorithms apply alternating exponentials of the cost and driving Hamiltonians. Indeed, when the QAOA depth p equals the number of Trotter steps M and the variational parameters are chosen according to a linear annealing schedule,

$$\gamma_m = s_m \cdot \Delta t = \frac{m}{p} \cdot \frac{t_{\text{sweep}}}{p}, \quad \beta_m = (1 - s_m) \cdot \Delta t = \left(1 - \frac{m}{p}\right) \cdot \frac{t_{\text{sweep}}}{p}, \quad (3.71)$$

the QAOA exactly recovers first-order Trotterized quantum annealing with total sweep duration t_{sweep} [338]. In this sense, the QAOA can be understood as a generalization of Trotterized annealing where the variational parameters γ_m and β_m are freed from the constraints imposed by a fixed adiabatic schedule. This connection has been formalized in several recent works [338, 341–343], which demonstrate that optimized QAOA parameters often converge toward smooth, annealing-like trajectories as the circuit depth increases.

However, the key distinction lies in the classical optimization freedom: while quantum annealing follows a predetermined schedule constrained by the requirements of the adiabatic theorem (Eq. 3.11), QAOA can discover parameter sequences that exploit diabatic transitions to circumvent small spectral gaps [341]. Numerical studies have shown that on problem instances where adiabatic quantum annealing fails due to exponentially small gaps, the QAOA can find high-quality solutions orders of magnitude faster by learning to utilize these non-adiabatic mechanisms [341].

Despite its promise, the QAOA faces several practical challenges. Chief among these is the barren plateau phenomenon [72], wherein the gradients of the cost function vanish exponentially with increasing system size for certain variational quantum circuits. While the original analysis focused on random parameterized circuits, subsequent work has examined problem-tailored ansätze like the QAOA using tools from quantum optimal control [344]. These studies show that the appearance of barren plateaus depends on the initial states and problem instances.

Several strategies have been proposed to mitigate these challenges. Initializing QAOA parameters based on Trotterized quantum annealing schedules has been shown to provide a favorable starting point that avoids many local minima in the optimization landscape [338].

Additionally, parameters optimized for small problem instances can often be transferred to larger instances, reducing the classical optimization overhead [341, 345]. These findings suggest that leveraging the connection to quantum annealing may be essential for practical implementations of the QAOA on near-term quantum devices.

Chapter 4

Noise mitigation for quantum annealing via dynamical decoupling

Realizing quantum annealing algorithms (Chapter 3) on near-term hardware requires protocols that remain effective in the presence of the noise affecting current quantum computing platforms. In trapped-ion implementations based on magnetic gradient induced coupling (MAGIC, see Chapter 1), fluctuations of the magnetic field and of the gradient itself translate into local field disorder acting on the ions. Typical amplitudes of these fluctuations can exceed the achievable two-body coupling strengths by more than one order of magnitude (see Sec. 1.1.3), well beyond the level at which the ground state of the cost Hamiltonian is no longer guaranteed to encode the solution of the underlying optimization problem. Mitigating local-field noise is therefore a crucial prerequisite for any practically useful quantum annealing protocol running on the MAGIC platform.

In this chapter we develop and analyze an active noise-mitigation scheme tailored to quantum annealing in trapped-ion systems. The scheme relies on periodically applied global spin-flip pulses, that is, the simplest form of dynamical decoupling, combined with an ancilla-qubit encoding that maps all problem-defining local fields onto two-body interactions. This construction renders the full annealing Hamiltonian invariant under the decoupling pulses, while the leading effect of the pulses is to cancel local-field disorder. We support this picture analytically through a Magnus expansion of the time-evolution operator and numerically through full simulations of annealing sweeps subject to time-dependent noise. The numerical simulations use experimentally motivated parameters for the MAGIC platform [163, 205] and a noise model dominated by the 50 Hz electrical-grid frequency and its odd harmonics, which is characteristic of unshielded laboratory environments. Across all considered settings, we find that modest dynamical decoupling pulse rates of about 2.5 pulses per millisecond, well within the reach of current trapped-ion experiments, are sufficient to recover the noise-free annealing fidelity even for noise amplitudes corresponding to a worst-case, unshielded MAGIC setup.

The remainder of this chapter is organized as follows. In Sec. 4.1, we quantify the impact of static disorder in the two-body couplings and in the local fields on the ground state of the cost Hamiltonian in a quantum annealing protocol, identifying local-field fluctuations as the most relevant error source under realistic experimental conditions. Section 4.2 introduces the

dynamical decoupling protocol, derives the corrections to the ideal annealing evolution via a Magnus expansion, and discusses how the same construction can be repurposed to modulate two-body coupling strengths. Section 4.3 presents numerical simulations of the full protocol for two industrially relevant benchmark problems, namely Multiple-Object Tracking and cutting stock, as well as the Sherrington–Kirkpatrick spin-glass model as an unstructured benchmark, and concludes with a universal scaling analysis connecting the results.

All content in this chapter is based on Ref. [144].

4.1 Effect of static disorder on quantum annealing performance

Before introducing the dynamical decoupling protocol, we first quantify the sensitivity of the annealing protocol to static disorder, that is, disorder terms that remain approximately constant on the timescale of a single annealing sweep. Two qualitatively different sources are relevant for the MAGIC platform (see Chapter 1). Fluctuations in the two-body couplings J_{ij} can be induced by slow drifts of the trapping frequency or, in Trotterized implementations, by imperfect two-body gate operations. Longitudinal-field fluctuations δh_i^z instead originate from magnetic-gradient variations and from stray external magnetic fields. For each source, we numerically determine the disorder amplitude beyond which the solution-encoding ground state of the cost Hamiltonian changes, using the minimal MOT instance with five qubits, constructed in Sec. 3.2.3, as a test problem.

4.1.1 Fluctuations in the two-body couplings




Figure 4.1: Probability of a changed ground state of the cost Hamiltonian for the quadratized minimal MOT problem (five qubits, see Eq. (3.44)) 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 $\leq 2.5 \cdot 10^{-5}$. The grey dashed line corresponds to a fit with an arctan function, starting at $\sigma(\delta J_{ij})/J = 0.3$. Fit parameters: $a, b, c, d \approx 0.41, 2.59, 0.50, 0.22$.

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

$$\omega_z \rightarrow \omega_z + \delta\omega_z, \quad (4.1)$$

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. (1.33), we calculate the two-body coupling variations to first order in $\delta\omega_z$ as

$$\delta J_{ij} \approx \sum_{n=1}^N \frac{-\hbar \partial_z \omega_i \partial_z \omega_j}{\nu_n^2 \omega_z m} S_{in} S_{jn} \delta\omega_z = -2 \frac{J_{ij}}{\omega_z} \delta\omega_z, \quad (4.2)$$

where we used the fact that the axial phonon modes ν_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. 3.1.1), 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 Sec. 4.3.3), 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 5 qubits, constructed in Sec. (3.2.3), as a test problem. Each data point in Fig. 4.1 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 $p(1-p)/10000 \leq 2.5 \cdot 10^{-5}$. 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 Chapter 1), 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 $\sim 99.9\%$ can already be achieved for current trapped-ion quantum computers [164]. Although we cannot exclude that the picture may change when scaling up the size of the optimization problem, the 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.

4.1.2 Local field fluctuations

Magnetic field gradient fluctuations in the MAGIC setup, 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. (3.8)). We model these disorder terms by modifying the local field components:

$$h_i^z \rightarrow h_i^z + \delta h_i^z, \quad (4.3)$$

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. 3.1.1). 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. 4.3.

To assess the quantum annealing protocol's resilience to local field variations, we add disordered diagonal terms δh_i^z to the cost Hamiltonian given in Eq. 3.44. 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.

Figure 4.2 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 characteristic energy scale of the physical platform, here set by the maximal two-body coupling strength (see Sec. 4.3.1). The grey dashed 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 [163], a five-ion chain reaches a maximal two-body coupling strength of $J_{ij} \approx 2\pi \times 26$ Hz [205]. In the same setup, local field fluctuations can reach magnitudes on the order of 10^3 rad/s in the worst-case scenario without any magnetic shielding




Figure 4.2: Probability of a changed ground state of the cost Hamiltonian for the quadratized minimal MOT problem (using 5 qubits, see Eq. 3.44) 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 $\leq 2.5 \cdot 10^{-5}$), where δh_i^z is sampled from a Gaussian distribution with zero mean. The grey dashed 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: $a, b, c, d \approx 0.57, 0.81, 0.99, 0.10$.

and stabilization [232], corresponding to $\delta h_i^z/J \sim 10$. 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 $\max_{ij} |J_{ij}| \approx 2\pi \times 1650.6$ Hz [205], for which the worst-case ratio $\delta h_i^z/J$ falls to about 0.1, below the empirical robustness threshold identified above. However, the ratio rapidly deteriorates for larger ion chains, where the spin-spin couplings decrease (see Fig. 1.1), and dedicated mitigation strategies are still required to keep the protocol robust as one scales up to industrially relevant problem sizes. The next section demonstrates that local fields can be strongly suppressed through simple dynamical decoupling sequences, even when they fluctuate substantially and when they are time dependent.

4.2 Applying dynamical decoupling pulses during annealing

We now turn to the construction and analysis of the dynamical decoupling protocol itself. We restrict the discussion to the simplest non-trivial pulse sequence, namely periodically applied global π -pulses, which act as instantaneous spin flips and can be realized in the MAGIC setup as resonant microwave drives of finite duration (see also discussion in Sec. 1.3.1). In Sec. 4.2.1, we show that these pulses, combined with the ancilla-qubit encoding introduced in Sec. 3.1.1, leave the annealing Hamiltonian invariant while periodically reversing the sign of any longitudinal-field disorder. Section 4.2.2 formalizes this picture through a Magnus expansion of the time-evolution operator over a single pulse interval, identifying the leading-order corrections to the ideal annealing evolution and estimating the resilience of the protocol to imperfect pulse timing. In Sec. 4.2.3, we extend the construction to spin flips acting on

subsets of qubits and discuss how the same mechanism can be used to effectively modulate selected two-body coupling strengths, providing a practical method for encoding QUBO problems in the native all-to-all couplings of the MAGIC platform.

4.2.1 Effect of dynamical decoupling on quantum annealing protocols

Dynamical decoupling represents a powerful error mitigation technique designed to extend the coherence time of quantum systems by suppressing environmental decoherence [119, 132, 135, 239, 346–350]. 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. (3.9). 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 towards the correct solution.

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. (1.3.1)). This choice of simple π -pulses is sufficient for mitigating the local field fluctuations considered in our analysis (see Sec. 4.1.2). The action of such pulses on the driver and cost Hamiltonians of the annealing protocol yields

$$e^{i\frac{\pi}{2}\sigma_x} H_{\text{driving}} e^{-i\frac{\pi}{2}\sigma_x} = H_{\text{driving}}, \quad (4.4)$$

$$e^{i\frac{\pi}{2}\sigma_x} H_{\text{cost}} e^{-i\frac{\pi}{2}\sigma_x} = \sum_{ij} J_{ij} \sigma_i^z \sigma_j^z - \sum_i h_i^z \sigma_i^z. \quad (4.5)$$

Notably, the spin-flip operations leave the driver 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. 3.1.1. By encoding all local fields into additional two-body interactions through the introduction of an ancilla qubit, see Eq. (3.6), 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 Sec. 4.3.3.

Practical experimental implementations often employ more sophisticated dynamical decoupling sequences. For instance, universally robust sequences [135, 136, 239] 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. [143, 351].

4.2.2 Analytical treatment via Magnus expansion

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 modified quantum annealing Hamiltonian discussed in Sec. 3.1.2 (constant cost Hamiltonian) with noise,

$$\begin{aligned} H(t) &= C(t)H_{\text{driving}} + H_{\text{cost}} + H_{\text{noise}} \\ &= -C(t) \sum_i h^x \sigma_i^x + \sum_{ij} J_{ij} \sigma_i^z \sigma_j^z + \sum_i \delta h_i^z(t) \sigma_i^z, \end{aligned} \quad (4.6)$$

where the problem is encoded only in the two-body interactions, 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. eq4.6 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

$$C(t) = C(t_-) - v(t_-)t, \quad t \in [t_-, t_+]. \quad (4.7)$$

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. Further, 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 [352]:

$$\begin{aligned} U_{t_- \rightarrow t_0} &= \mathcal{T} e^{-i \int_{t_-}^{t_0} dt' H(t')} \\ &= e^{-i \int_{t_-}^{t_0} dt' H(t') - \frac{1}{2} \int_{t_-}^{t_0} dt' \int_{t_-}^{t'} dt'' [H(t'), H(t'')] + \mathcal{O}(\Delta t^4)}. \end{aligned} \quad (4.8)$$

Evaluating the commutators, integrating $C(t)$ (using Eq. (4.7)) and keeping only terms up to third order in Δt gives

$$\begin{aligned} U_{t_- \rightarrow t_0} &\approx \exp(-i[C(\bar{t}_-)H_{\text{driving}} + H_{\text{cost}} + H_{\text{noise}}]\Delta t \\ &\quad - \frac{i}{6} h^x v(t_-) \left[\sum_i \delta h_i^z \sigma_i^y + 2 \sum_{ij} J_{ij} \sigma_i^z \sigma_j^y \right] \Delta t^3), \end{aligned} \quad (4.9)$$

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),

$$\sigma^x U_{t_0 \rightarrow t_+} \sigma^x U_{t_- \rightarrow t_0} = e^{-i[\hat{A}_1 + \hat{A}_2 + \hat{A}_3]2\Delta t + \mathcal{O}(\Delta t^4)}, \quad (4.10)$$

with

$$\hat{A}_1 = \sum_{ij} J_{ij} \sigma_i^z \sigma_j^z - h^x C(t_0) \sum_i \sigma_i^x, \quad (4.11)$$

$$\hat{A}_2 = \left(h^x C(t_0) \sum_i \delta h_i^z \sigma_i^y \right) \Delta t, \quad (4.12)$$

$$\hat{A}_3 = \left(-\frac{2}{3} h^x v \sum_{ij} J_{ij} \sigma_i^z \sigma_j^y + \frac{2}{3} h^x C(t_0) \sum_i \delta h_i^{z^2} \sigma_i^x \right) \Delta t^2. \quad (4.13)$$

Here, we write $\sigma^x = \bigotimes_i \sigma_i^x$ as shorthand notation denoting the global spin flip due to the dynamical decoupling pulses. Further, 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. (4.12), 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. Further, in Section 4.3.4, 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. (4.13), are composed of two parts: First, there is a two-body interaction term whose strength is proportional to $\sim h^x v J_{ij} \Delta t^2$. 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 ($\sim \delta h_i^{z^2} \Delta t^2$).

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. [136]. 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 [264–266].

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. (4.6)) 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. Afterwards 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

$$\sigma^x U_{t_0^\epsilon \rightarrow t_+} \sigma^x U_{t_- \rightarrow t_0^\epsilon} = e^{-i\hat{A}_1^\epsilon 2\Delta t + \mathcal{O}(\Delta t^2)}, \quad (4.14)$$

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:

$$\hat{A}_1^\epsilon = H_{\text{cost}} + C(t_0)H_{\text{driving}} + \frac{\epsilon}{\Delta t}H_{\text{noise}}. \quad (4.15)$$

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. 4.1.2 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 (see Sec. 4.3.1 for a discussion of the energy scales). In our numerical simulations of annealing sweeps with time-dependent local-field noise in Sec. 4.3, we observe that moderately larger fluctuations in pulse spacing of the order of 2% also have negligible impact on performance, provided they do not exhibit a systematic pattern that would enable coherent error accumulation over the course of the annealing sweep.

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 $\delta h_i^z \Delta t$.

4.2.3 Modulating coupling strengths with dynamical decoupling

In Sec. 4.2, 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. 4.6). In this section 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 [236, 353]. This scheme can be potentially utilized to encode QUBO problems into the native all-to-all couplings of the MAGIC setup (see Sec. 1.2.1).

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. 4.2, 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 Hamiltonian, 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

$$U_{t_- \rightarrow t_0} \approx \exp\left(-i \sum_{ij} J_{ij} \sigma_i^z \sigma_j^z \Delta t_- - iC(\bar{t}_-)H_{\text{driving}}\Delta t_-\right), \quad (4.16)$$

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. (4.16), and the evolution over the entire time interval can be calculated to first order as

$$\begin{aligned} \sigma_k^x U_{t_0 \rightarrow t_+} \sigma_k^x U_{t_- \rightarrow t_0} \approx \exp \bigg(& -i \sum_{i,j \neq k} J_{ij} \sigma_i^z \sigma_j^z \Delta t \\ & -i \sum_{i \neq k} J_{ik} \sigma_i^z \sigma_k^z (\Delta t_- - \Delta t_+) \\ & -i C(\bar{t}) H_{\text{drive}} \Delta t \bigg), \end{aligned} \quad (4.17)$$

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 Sec. 4.3.3), 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. 4.2. 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. (4.17), 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.

4.3 Numerical simulations of dynamical decoupling protocol

We now turn to numerical simulations of the dynamical decoupling protocol introduced in the last section for fully time-dependent local-field noise. Throughout this section, we use the minimal MOT instance with five qubits as our primary benchmark, with additional results for a nine-qubit MOT encoding, cutting stock instances of five and six qubits, and Sherrington–Kirkpatrick spin-glass instances of up to twelve qubits, allowing us to test the protocol across problems of different structure and size. Section 4.3.1 fixes the energy scales and unit conventions used throughout. Section 4.3.2 specifies the noise spectra and numerical details how individual noise realizations are generated, as well as how the dynamical decoupling pulses are distributed across the discretized annealing sweep. Section 4.3.3 presents the simulation results for the considered benchmark problems, including a comparison with an alternative ancilla-free encoding. Finally, Sec. 4.3.4 demonstrates that the fidelity curves obtained for different noise amplitudes collapse onto a single curve when expressed as a function of a generalized scaling parameter, in agreement with the analytical predictions of Sec. 4.2.2.

4.3.1 Energy scales and unit conventions

In the remainder of this section, we report physical time in units of the inverse characteristic energy scale of the cost Hamiltonian (see Sec. 3.1.4). Since the all-to-all couplings J_{ij} in the MAGIC scheme are not uniform across the chain (see Sec. 1.2.2), we define this scale as

$$J \equiv \frac{1}{2} \max_{ij} |J_{ij}|, \quad (4.18)$$

i.e., as half the strongest two-body coupling in the specific experimental setup (the factor of $1/2$ is due to convention in defining the MAGIC Hamiltonian, Eq. (1.32)). All other quantities with the dimension of an energy or an angular frequency, including local-field disorder amplitudes, Rabi frequencies, and dynamical decoupling pulse rates, are quoted as angular frequencies, with the convention that values stated in Hz are to be read as angular frequencies in rad/s unless an explicit prefactor of 2π is given (in which case the unit is just 1/s).

The relevant numerical values for the MAGIC platform considered in this work were established in Sec. 1.2.2. For a chain of five $^{171}\text{Yb}^+$ ions in an axial trap with $\omega_z = 2\pi \times 130$ kHz, the experimentally demonstrated magnetic field gradient $\partial_z B = 19 \text{ T m}^{-1}$ [163] yields $\max_{ij} |J_{ij}| \approx 2\pi \times 26.5$ Hz, while the near-future value $\partial_z B = 150 \text{ T m}^{-1}$, accessible with current technology [238, 247], increases this to $\max_{ij} |J_{ij}| \approx 2\pi \times 1650.6$ Hz. Through Eq. (4.18), these correspond to characteristic energy scales of $J \approx 2\pi \times 13.3$ Hz and $J \approx 2\pi \times 825$ Hz, respectively. Local-field fluctuations in the same setup can reach amplitudes on the order of 10^3 rad s^{-1} in the worst case, without magnetic shielding, but can be reduced through active stabilization [232].

For the MOT and Sherrington–Kirkpatrick benchmarks discussed below, we adopt the more conservative value $J = 26 \text{ rad s}^{-1}$, corresponding to a gradient of $\partial_z B \approx 10.6 \text{ T m}^{-1}$, i.e., somewhat below the value already realized in Ref. [163]. With the dimensionless annealing time of $2.6 J^{-1}$ (optimized for large success probabilities after the anneal), this fixes the physical sweep duration to 100 ms. The cutting stock instances exhibit substantially smaller minimal spectral gaps along the sweep and consequently require a longer dimensionless annealing time of $26 J^{-1}$ for comparable noiseless fidelities; to keep the physical sweep duration at 100 ms and enable better comparison to the other benchmarks, we assume $J = 260 \text{ rad s}^{-1}$ for those instances. Since all numerical results are expressed in dimensionless units of J , they remain unchanged for any other choice of gradient. As reference points, the experimentally demonstrated gradient $\partial_z B = 19 \text{ T m}^{-1}$ would shorten the physical sweep of the MOT problem to approximately 30 ms, and the near-future value $\partial_z B = 150 \text{ T m}^{-1}$ to approximately 0.5 ms.

4.3.2 Generating time-dependent local-field noise

In this section, we explain how the noise spectra that we use in the numerical simulations of the following section are defined and local field fluctuation samples are generated. Further, we discuss how the dynamical decoupling pulses are distributed throughout the numerically discretized annealing sweep.

Generating the noise samples — As discussed in Sec. 1.1.3, 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:




Figure 4.3: Example of a time trace of local field noise with 50000 sampling points over 100 milliseconds, generated from a noise spectrum with two Lorentzian peaks at 50 Hz and 150 Hz (Eq. 4.19) with a noise amplitude of 750 Hz. The inset shows a zoom onto the first 0.2 milliseconds (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. 4.2).

$$S(f) = \frac{1}{2\pi} \left(\frac{\gamma}{(f - f_1)^2 + \gamma^2} + \frac{\beta}{(f - f_3)^2 + \beta^2} \right), \quad (4.19)$$

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 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. [354]. The sampling covers $N = 50\,000$ points over 100 ms, corresponding to our in this section adopted 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. 4.5 in the next section 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. 4.5, are generated analogously to the method outlined above.

In Fig. 4.3, 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 Hz 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 milliseconds (corresponding to the time of an annealing sweep used in our numerical simulations below), while the inset zooms in on the first 0.2 milliseconds. This

smaller time interval corresponds to the time between two dynamical decoupling pulses for the highest pulse rate of 5 kHz we consider (see Sec. 4.3.3) and is resolved into 100 distinct noise sampling points.

Numerical simulation timescales — Our numerical implementation involves two distinct timescales that require careful 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. 4.2.2). 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.

4.3.3 Numerical results for MOT, cutting stock and SK model

In this section, we present numerical simulations of the dynamical decoupling protocol for various benchmark problems. We employ the minimal MOT instance with five qubits, defined in Eq. (3.44) as our first test problem in Sec. 4.3.3, after which we extend the analysis to the 9-qubit instance constructed in Sec. 3.2.3. Additionally, in Sec. 4.3.3 and Sec. 4.3.3 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. Finally, in Sec. 4.3.3, we discuss an alternative annealing protocol without ancilla qubit solving the first minimal MOT problem instance using only four qubits.

Multiple-Object Tracking

The first benchmark test case is the industrially-relevant MOT computer vision problem. We consider two specific problem instances, one using 5 qubits and one using 9 qubits, both constructed in Sec. 3.2.3. Both problems are encoded using only 2-body interactions (see Sec. 3.1.1 for a discussion on the encoding method). Furthermore, we consider the modified annealing schedule of Eq. (3.9) for all simulations, where the cost Hamiltonian is held constant during the annealing evolution.




Figure 4.4: 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 Hz 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 Section 4.3.2). 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.

5-qubit MOT instance Our simulations solve the time-dependent Schrödinger equation using the annealing Hamiltonian with 5 qubits defined in Eq. (4.6), where the 2-body couplings are defined by the 5-qubit MOT problem instance constructed in Sec. 3.2.3. 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, we adopt the conservative energy scale $J = 26$ Hz which is below values demonstrated in the MAGIC setup [163] and corresponds to an annealing time of 100 ms for our chosen parameters. We set the driving strength to $h_x = 3J$ 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. 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




Figure 4.5: Impact of noise peak frequency on dynamical-decoupling performance. As Fig. 4.4, but using different noise spectra with a single Lorentzian peak centered at varying frequencies, all with the same amplitude of 750Hz. 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 750Hz. The horizontal dashed line shows the achievable fidelity of the ideal noiseless protocol. Data for each boxplot corresponds to 50 random independent noise realizations (see Section 4.3.2). Similar to Fig. 4.4, 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.

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. 1.1.3, 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 discussed in the previous section. At approximately equidistant intervals across the 50 000 discrete time steps, 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 two-fold degenerate true ground state encoding the problem solution. Figure 4.4 presents this protocol fidelity as a function of the number of DD pulses for a noise spectrum featuring two Lorentzian peaks centered at 50 Hz 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 [232].

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 grey line at $\approx 84\%$ represents 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 towards 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 $2\pi \times 100$ kHz. When setting the Rabi frequency to $2\pi \times 100$ kHz, the upper limit for the pulse rate would be around 100 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. 4.4, marked by 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. 4.4, with convergence to ideal fidelity occurring slightly faster.

Figure 4.5 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. 4.4) 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. 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 [238, 247], enabling coupling strengths up to $J \approx 2\pi \times 1650$ Hz [205]. In such scenarios, the annealing time required to achieve fidelities comparable to those in our simulations would reduce from 100 ms

to approximately 0.5 ms. Moreover, the ratio of noise amplitude to coupling strength would become more favorable, enabling high-fidelity sweeps with even fewer dynamical decoupling pulses.




Figure 4.6: 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, using the MOT test problem with 9 qubits (see Sec. 3.2.3) and an annealing time of $2.6/J$ (corresponding to a sweep duration of 100ms for an energy scale of $J = 26$ Hz). Box plots compare results for correlated noise using identical noise spectra with two Lorentzian peaks at 50 Hz and 150 Hz (see Sec. 4.3.2) 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 gray dashed line shows the achievable fidelity of the ideal noiseless protocol (≈ 0.74). Data for each boxplot corresponds to 50 random independent noise realizations (see Sec. 4.3.2). 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.

9-qubit MOT instance Figure 4.6 presents our numerical results for the 9-qubit problem under the same correlated local field noise conditions discussed in Sec. 4.2. The time-dependent noise is sampled from a spectrum with two dominant frequency peaks at 50 Hz and 150 Hz, analogous with the 5-qubit analysis (Fig. 4.4). We assume the same energy scale of $J = 26$ Hz to enable direct comparison with the results for the 5-qubit problem, 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 [205].

The 9-qubit problem behaves similarly to its 5-qubit counterpart. Final fidelities converge

toward the ideal noiseless performance (gray dashed 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 the specific problem instance and system size (further supported by our results on the cutting stock problem and the Sherrington-Kirkpatrick model in the following sections). 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. 4.2.1, where corrections to the ideal annealing protocol are shown to be independent of the specific coupling strengths that encode the optimization problem.

Cutting stock problem

The second benchmark problem we consider is the cutting stock problem, which is also industrially relevant, similar to the MOT problem. We discussed this problem in Sec. 3.2.4, where we also constructed specific minimal problem instances with 5 and 6 qubits respectively (see Fig. 3.5). Note that we again consider the cost Hamiltonian encoding without local fields, which uses an additional ancilla qubit (see Sec. (3.1.1)). Figure 4.7 presents numerical quantum annealing results for cutting stock instances with 5 and 6 qubits, and the MOT instances with 5 and 9 qubits from the previous section, all subject to identical correlated local-field noise. The time-dependent noise is sampled from a spectrum characterized by two dominant Lorentzian peaks at 50 Hz 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. 4.4 and 4.6.

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. [317] 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 (see Sec. 4.3.1). 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. 4.7(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 4.7(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.

Sherrington-Kirkpatrick model

While we focused above on application-motivated optimization problems (MOT and cutting stock), here we demonstrate that dynamical decoupling (DD) equally mitigates local-field




Figure 4.7: Annealing results for MOT instances with five and nine qubits (see Sec. 3.2.3), and for cutting stock instances with five and six qubits (see Fig. 3.5), 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. 4.4 and 4.6, 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.

noise in problems without internal structure. As a representative example, we consider the Sherrington–Kirkpatrick (SK) spin-glass model with local fields [355, 356], 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

$$H_{\text{SK}} = \sum_{i \neq j} J_{ij}^{\text{SK}} \sigma_i^z \sigma_j^z + \sum_i h_i^{\text{SK}} \sigma_i^z, \quad (4.20)$$

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 above, local fields are replaced by two-body terms through the introduction of a single ancilla qubit, ensuring that all problem




Figure 4.8: 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. 4.4 and 4.6, defined by Eq. (4.19), 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.

information is contained in the interaction terms only (see Sec. 3.1.1). This allows local-field noise to be mitigated via DD without altering the encoded cost function.

Figure 4.8 summarizes our numerical results. Analogous to the results in the previous sections, we consider a noise spectrum with two Lorentzian frequency peaks at 50 Hz and 150 Hz (Eq. (4.19)), and a noise amplitude of 750 Hz. As for the other benchmark problems, we consider the unit-less 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 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.

This observation also matches our analytical treatment in Sec. 4.2, 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.

Problem encoding without ancilla qubit

In the previous discussions, we considered 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 Sec. 3.1.1). 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. 4.2). 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 desirable to retain the original QUBO formulation that directly incorporates local field terms in the cost Hamiltonian. In this section, 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

$$H = C(t)H_{\text{driving}} + \sum_{ij} J_{ij}\sigma_i^z\sigma_j^z + \sum_i (h_i^z + \delta h_i^z(t))\sigma_i^z. \quad (4.21)$$

Implementing the same dynamical decoupling scheme as in Sec. 4.2 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 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. E.g., when the problem-encoding local fields are implemented via detuning of a microwave signal from resonance of the corresponding ion transition (see for example discussion in Ref. [205]), it suffices to quickly switch the sign of the detuning simultaneously with the action of the dynamical decoupling pulse. Other implementations rely on virtual Z rotations [253] 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 previous sections when applied to the minimal MOT problem (see Sec. 4.3.3). 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 4.9 presents the final noiseless fidelities achieved for both protocols across varying annealing sweep durations, expressed in units of the inverse characteristic energy scale J . We consider again a sweep duration of $2.6J^{-1}$ for numerical calculations, corresponding to 100 ms when assuming $J = 26$ Hz (indicated by the vertical red dashed 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.




Figure 4.9: Final noiseless fidelities achieved for different total linear annealing sweep durations for the minimal MOT QUBO problem (Sec. 3.2.3). We compare the performance of the annealing protocol considered in the previous sections (Eq. 3.9), encoding the problem exclusively into the two-body interactions using 5 qubits overall (orange), versus the direct encoding utilizing local fields with 4 qubits overall (blue). The annealing time is given in units of the inverse characteristic energy scale J . The vertical red dashed line denotes the annealing time considered for the numerical simulations of the MOT problem in Sec. 4.3.3 ($2.6J^{-1}$).

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. 4.4). 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.

4.3.4 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. 4.4, 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.

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. 4.2.1. There, the leading-order corrections due to local field fluctuations appear as perturbations to the driver Hamiltonian proportional to $\delta h_i^z \times \Delta t$, see Eq. (4.12), 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. 4.4 and 4.5). This implies that the deviation from the ideal (noiseless) annealing performance should be governed by this product.




Figure 4.10: 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. 4.2.1). The grey 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 Hz 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 test this prediction, we plot in Fig. 4.10 (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. 4.10 (a) shows an exponential fit to the combined data.

However, the perturbative argument in Sec. 4.2.1 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. 4.10(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 Hz 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.

Taken together, the analytical and numerical results of this chapter establish that simple periodic π -pulses suffice to render quantum annealing on the MAGIC platform robust against the dominant longitudinal-field noise, at pulse rates well within the reach of current trapped-ion experiments and largely independent of the specific cost Hamiltonian and system size. This removes one of the main practical obstacles to running industry-relevant optimization problems on near-term MAGIC hardware, and quantum computing platforms in general, and clears the way for the performance-oriented improvements of the annealing protocol itself, which we turn to in the following chapter.

Chapter 5

Improving quantum annealing performance

In Chapter 3 we discussed how the time required to solve an optimization problem with quantum annealing is, in the adiabatic regime, controlled by the inverse square of the minimum energy gap ΔE encountered along the sweep (Sec. 3.1.3). For hard combinatorial problems, this gap typically closes exponentially with the problem size, $\Delta E \sim e^{-\alpha N}$, a hallmark of a first-order quantum phase transition along the annealing path. The adiabaticity time then grows as $T \propto e^{2\alpha N}$, which is the central bottleneck limiting the practical applicability of the standard linear annealing protocol with a transverse-field driver.

Considerable effort has therefore been devoted to designing modified annealing protocols whose minimum gap closes more slowly with system size, or whose performance benefits from controlled non-adiabatic dynamics. In this chapter, we discuss three promising avenues toward this goal, each acting on a different component of the time-dependent annealing Hamiltonian.

The main focus of the chapter, presented in Sec. 5.1, concerns the cost Hamiltonian itself. Combinatorial optimization problems are conventionally cast in QUBO form before being implemented on quantum annealing hardware, even when their natural formulation involves higher-order interactions. We show that retaining the native PUBO structure of the cost Hamiltonian can yield significantly larger minimum gaps relative to the characteristic energy scale of the device and, for certain instance families, an exponentially improved gap-scaling exponent. Using the paradigmatic 3-SAT problem as a benchmark, we quantify the resulting savings in qubit count and adiabaticity time, and show that the gain persists when accounting for the overhead of realizing higher-order interactions on actual hardware.

Section 5.2 discusses modifications of the driving Hamiltonian, reviewing the strategy of inhomogeneous quantum annealing, in which different qubits are driven across the critical phase transition region at different times. We summarize the main proposals from the literature and the resulting reports of polynomial and, in idealized mean-field limits, exponential improvements over homogeneous protocols.

Finally, Sec. 5.3 introduces a third avenue, the augmentation of the annealing path by a non-stoquastic catalyst term that is active only during the intermediate stages of the sweep. We review the notion of stoquasticity and its connection to the sign problem of quantum Monte Carlo, and discuss the cases in which non-stoquastic catalysts soften first-order transitions or enable diabatic shortcuts to the solution. We will revisit non-stoquastic annealing in Chapter 6, where we analyze quantum resources generated along such paths, in

particular bipartite entanglement and nonstabilizerness of the instantaneous ground state.

The content of this chapter is based on Refs. [203] and [204].

5.1 Utilizing higher-order interactions in the cost Hamiltonian

As discussed in Chapter 3, to solve combinatorial optimization problems with quantum annealing, they have to be encoded into a cost Hamiltonian whose ground state represents the desired solution (Sec. 3.1.1). When the underlying classical cost function is at most quadratic in the binary variables, this Hamiltonian contains only local fields and two-body interactions, and the optimization problem is of QUBO type. Many problems of practical interest, however, are naturally formulated as PUBO problems of higher degree, the canonical example being the NP-complete 3-SAT problem, which is cubic in its native form (Sec. 3.2.5). Since most currently available quantum annealing hardware supports only local fields and two-body interactions, the standard approach is to reduce any PUBO instance to an equivalent QUBO problem by introducing ancillary slack variables together with energy penalties enforcing their consistency. This reduction incurs a twofold cost: each slack variable requires an additional qubit, and the penalty terms must dominate the problem-encoding interactions, which under the limited coupling rates of physical hardware reduces the effective energy scale available to the original problem and, with it, the achievable sweep speed.

In this section, we investigate whether retaining the native PUBO structure of the cost Hamiltonian, rather than reducing it to QUBO, can improve the performance of quantum annealing. The relevant figure of merit is the minimum energy gap encountered along the annealing sweep, since for hard problems its closing controls the adiabaticity time and thus the time-to-solution (Sec. 3.1.3). Using 3-SAT as a paradigmatic benchmark, we find that the native PUBO formulation yields larger minimum gaps relative to the characteristic energy scale and, for certain instance families, an exponentially improved gap-scaling exponent. These advantages persist when accounting for the overhead of realizing higher-order interactions on the underlying quantum hardware, both on analog and on digital platforms.

The remainder of this section is organized as follows. Sections 5.1.1 and 5.1.2 review the reduction of a PUBO problem to QUBO via slack variables and discuss the resulting hierarchy of physical energy scales, including the role of the penalty strength. Sections 5.1.3 and 5.1.4 present our numerical analysis of the minimum gap for 3-SAT instances generated by three different generators, culminating in the system-size scaling of the minimum spectral gap for both formulations. Sections 5.1.5 and 5.1.6 address the implementation of three-body interactions on quantum hardware and combine the gap-scaling results with the resulting overhead to quantify the speedup of PUBO over QUBO.

5.1.1 Reducing PUBO to QUBO

The reduction of a PUBO cost function to QUBO form rests on a single elementary substitution, already introduced in Sec. 3.2.5 in the context of 3-SAT. Any cubic term $x_i x_j x_k$ with binary variables in the cost function can be replaced by a quadratic term $x_i y$ involving an ancilla variable $y \in \{0, 1\}$, provided that the constraint $y = x_j x_k$ is enforced energetically through the standard penalty

$$P(x_j, x_k, y) = 3y + x_j x_k - 2x_j y - 2x_k y, \quad (5.1)$$

which evaluates to zero precisely when $y = x_j x_k$ and is strictly positive otherwise. Iterating this substitution removes all three-body terms from the cost function at the price of one ancilla qubit and one penalty contribution per substitution. PUBO problems of higher degree are reduced analogously by recursive quadratization [330, 357].

Two practical caveats accompany the procedure. First, finding the substitution that minimizes the number of introduced ancillas is itself NP-hard [357]. In our numerical analysis we employ the simple brute-force heuristic discussed in Sec. 3.2.5, in which the pair of variables that appears most frequently among the cubic terms is replaced first, followed by the next-most-frequent pair, and so on. Second, the penalty strength must be large enough that violating the constraint $y = x_j x_k$ is never energetically favorable, but not so large that it dominates the problem encoding terms. Since the maximum two-body coupling supported by a quantum annealing device is bounded by hardware, an excessively strong penalty effectively rescales the encoding part of the Hamiltonian to a small fraction of the available energy scale and slows down the protocol. The interplay between penalty strength, characteristic energy scale and minimum spectral gap is analysed quantitatively in Sec. 5.1.2.

Toy model

To illustrate the procedure on a minimal example, consider the cubic polynomial

$$f(x) = x^3 + x, \quad (5.2)$$

which we wish to minimize on the half-open interval $[0, 2^2)$. We use a small resolution with $n = 2$ integer bits and $m = 1$ fractional bit and approximate the continuous variable as

$$x = 2x_1 + x_2 + \frac{1}{2}x_3, \quad x_i \in \{0, 1\}. \quad (5.3)$$

Substituting Eq. (5.3) into Eq. (5.2) and simplifying with the identity $x_i^2 = x_i$ yields the third-order PUBO cost function

$$\begin{aligned} f(x_1, x_2, x_3) = & 10x_1 + 2x_2 + \frac{5}{8}x_3 \\ & + 18x_1x_2 + \frac{15}{2}x_1x_3 + \frac{9}{4}x_2x_3 \\ & + 6x_1x_2x_3, \end{aligned} \quad (5.4)$$

which contains a single cubic monomial $\propto x_1x_2x_3$ and attains its unique minimum at $(x_1, x_2, x_3) = (0, 0, 0)$, corresponding to $x = 0$.

To reduce Eq. (5.4) to QUBO form, we introduce a single ancilla $y = x_2x_3$ and replace the cubic term by $6x_1y$. The resulting QUBO cost function is

$$f_{\text{QUBO}}(x_1, x_2, x_3, y) = f_{\text{quad}}(x_1, x_2, x_3) + 6x_1y + \mu P(x_2, x_3, y), \quad (5.5)$$

where f_{quad} collects the linear and quadratic terms of Eq. (5.4), P is the penalty defined in Eq. (5.1), and $\mu > 0$ controls its strength. For any sufficiently large μ , the global minimum of Eq. (5.5) coincides with that of Eq. (5.4) with a consistent $y = x_2x_3$. The QUBO formulation therefore reproduces the original optimization problem at the price of one additional binary variable and one penalty contribution.

This minimal example shows the two qualitative trade-offs of the QUBO reduction. When translating the classical cost functions into a cost Hamiltonian (see Sec. 3.1.1), the PUBO formulation requires only three qubits but contains a three-body interaction whose physical implementation is more demanding than two-body couplings (see Sec. 2.3). The

QUBO reduction removes the cubic term but adds a fourth qubit and a penalty contribution that must dominate the encoding interactions. Whether the gain in interaction order outweighs the qubit and penalty overhead depends on the structure of the problem and on the details of the hardware platform, a question we address quantitatively in the following sections using 3-SAT instances as benchmarks.

3-SAT

The benchmark problem used throughout this chapter is the paradigmatic NP-complete 3-SAT problem, defined in Sec. 3.2.5 as the conjunction of M disjunctive clauses over N Boolean variables, each clause depending on exactly three literals. As detailed in Sec. 3.2.5, the natural cost Hamiltonian assigns to every clause an energy term that vanishes on satisfying assignments and equals one otherwise. After the substitution $x_i \rightarrow \frac{1}{2} - \hat{s}_i^z$ (see Sec. 3.1.1), this produces a three-local Hamiltonian with up to M distinct three-body terms, together with two-body and local contributions of comparable magnitude. The PUBO encoding therefore uses exactly N qubits, independently of the number of clauses.

To probe how the underlying instance structure affects the PUBO-versus-QUBO comparison, we adapt three different generators from Ref. [328], denoted **toughSAT**, **uniquePT1** and **uniquePT4**. They differ in the way clauses are constructed and, as a consequence, in the structure of the resulting cost Hamiltonian and in the number of ancillas required for a QUBO reduction.

toughSAT : Each clause is generated by selecting three of the N variables uniformly at random and negating each of them with probability $1/2$. The clauses are mutually independent, and we set the clause-to-variable ratio close to the critical value $(M/N)_{\text{crit}} \approx 4.24$ [326] in order to maximise the probability of generating computationally hard instances. We post-select the resulting random instances for those with a unique solution, since hard instances with unique ground state are the relevant test cases for our gap-scaling analysis. Because the clauses are independent random samples of triples of variables, each clause typically contributes a distinct cubic monomial to the cost function, and the cost Hamiltonian contains up to $M \approx 4.24 N$ unique three-body interactions. The QUBO reduction therefore requires up to M ancilla qubits, with smaller counts achievable when the brute-force quadratization can reuse the same ancilla for several clauses sharing a common pair of variables. Among the three generators considered here, **toughSAT** yields the largest ancilla overhead in the QUBO reduction.

uniquePT1 : Following Ref. [328], the generator first samples a target solution $s \in \{0, 1\}^N$ uniformly at random and then constructs $M = N + 6$ clauses iteratively. A small initial set of clauses pins the value of the first variable, while each subsequent clause fixes the value of one further variable given the values of two already-determined variables, thereby deterministically forcing s to be the unique solution. Despite the lower clause count, the resulting cost Hamiltonian still contains generic three-body interactions of magnitude comparable to its two-body and local contributions, and the minimum spectral gap along the standard annealing path closes exponentially in N (Sec. 5.1.4). The QUBO reduction requires up to $M = N + 6$ ancillas, significantly fewer than for **toughSAT** instances at the same N , while still yielding hard problem instances.

uniquePT4 : The same target solution s is sampled, but the generator now produces $M = 4N$ clauses by attaching, to every variable, four specifically chosen clauses involving that variable together with two further variables. The four clauses associated with a given variable are constructed so that, when summed in the cost function, all three-body and two-body contributions cancel exactly, leaving only a local field that pins the variable to

its target value s_i . The cost Hamiltonian therefore reduces to a sum of independent local fields, $\hat{H}_{\text{cost}} = -\sum_i h_i^z \hat{s}_i^z$, with a trivial product ground state that can be read off directly. No QUBO reduction is needed, no ancillas are introduced, and the optimization problem is solvable in time linear in N . We include `uniquePT4` here only as an example which, even though initially containing three- and two-body interaction terms, is actually trivial to solve (as reflected in the annealing energy spectrum, where the gap is not closing). Otherwise we restrict our PUBO-versus-QUBO comparison in the following sections to instances generated by `toughSAT` and `uniquePT1`.

5.1.2 Energy scales and penalty strength

The general framework of factoring out a global energy scale J from the annealing Hamiltonian was introduced in Sec. 3.1.4. We briefly recall the relevant convention here. For a PUBO problem of degree k_{max} , J is identified with the maximum physical rate at which the slowest, that is, the highest-order, interaction can be implemented, since on essentially all current hardware the rates of k -body interactions decrease with k and the highest-order term sets the bottleneck for the achievable sweep speed. The dimensionless cost Hamiltonian is normalized so that $\max |J^{(k_{\text{max}})}| = 1$. For the third-order PUBO formulation of 3-SAT we therefore use $J = \max_{ijk} |J_{ijk}^{(3)}|$, while for the QUBO reduction we use $J = \max_{ij} |J_{ij}^{(2)}|$. Local fields are assumed to be implementable at sufficiently high strengths that they pose no additional restriction. The strength of the driving Hamiltonian is fixed at $\hbar^x = J$ for the standard linear protocol of Eq. (3.8), in line with the analysis of Sec. 3.1.5.

The QUBO reduction of a native PUBO problem introduces structure that is absent in the PUBO formulation and that complicates the simple energy-scale picture above. The cost Hamiltonian contains, in addition to the part encoding the problem, the penalty terms required to enforce the consistency of the ancilla variables (Sec. 5.1.1). Decomposing $\hat{H}_{\text{cost}}$ into the two contributions,

$$\hat{H}_{\text{cost}} = \hat{H}_{\text{problem}} + \lambda \hat{H}_{\text{constraints}}, \quad (5.6)$$

the two-body coupling matrix splits accordingly,

$$J_{ij}^{(2)}(\lambda) = [J_{ij}^{(2)}]_{\text{problem}} + \lambda [J_{ij}^{(2)}]_{\text{constraints}}, \quad (5.7)$$

where the prefactor $\lambda \geq 0$ globally rescales the constraint terms while preserving their internal relative weights. The cost Hamiltonian is normalized for each value of λ to the maximum magnitude of the resulting two-body couplings, $J = \max_{ij} |J_{ij}^{(2)}(\lambda)|$. The default heuristic introduced in Sec. 3.2.5, which sets the prefactor of every penalty term to $|a_{ijk}| + 1$ for the cubic interaction $a_{ijk} x_i x_j x_k$ it removes, corresponds to $\lambda = 1$ in this notation (i.e., the penalty terms have different strengths individually, and the global scale is set to one).

The key implication of the decomposition in Eq. (5.6) is that the constraint terms must, in general, dominate the encoding terms in order to guarantee that the consistency conditions are not violated in the ground state. After normalization, this implies $\lambda |[J_{ij}^{(2)}]_{\text{constraints}}|/J \sim 1$ together with $|[J_{ij}^{(2)}]_{\text{problem}}|/J \ll 1$, so that the part of the Hamiltonian that actually encodes the optimization problem occupies only a small fraction of the feasible coupling strengths, set by the hardware. This reduction of the encoding interaction strengths is absent in the native PUBO formulation, where every coupling encodes the problem itself. For the third-order PUBO formulation of typical 3-SAT instances encountered in our numerics, the maximum ratio of three-body to two-body couplings is of order unity, so that no analogous dilution arises in the PUBO case; this conclusion may differ for larger instances or for problems with a different PUBO structure.

Optimal penalty strength. We now examine quantitatively how varying the global prefactor λ in Eq. (5.6) affects the minimum spectral gap and, through it, the adiabaticity time of the annealing protocol. As a representative example, Fig. 5.1a shows the minimum gap along the annealing sweep of the QUBO reduction of a random 3-SAT instance generated by `uniquePT1` with $N = 6$ variables, as a function of λ . The dependence is non-monotonic, with a maximum at $\lambda \approx 0.64$ where $\Delta E/J \approx 0.074$, while the default choice $\lambda = 1$ yields $\Delta E/J \approx 0.06$, only about 19 % below the optimum. The behavior on either side of the optimum can be understood directly from the low-energy spectrum of $\hat{H}_{\text{cost}}$, plotted as a function of λ in Fig. 5.1b, with constraint-satisfying states shown in green and constraint-violating states in red.




Figure 5.1: Dependence of the minimum gap on the penalty strength in the QUBO reduction of a random 3-SAT instance generated by `uniquePT1` with $N = 6$ variables. (a) Minimum gap $\Delta E/J$ along the annealing sweep versus the penalty strength λ in Eq. (5.6). The gray dashed line marks the default choice $\lambda = 1$, for which $\Delta E/J \approx 0.06$; the optimum, $\Delta E/J \approx 0.074$, occurs at $\lambda \approx 0.64$. (b) Lowest 30 eigenvalues of $\hat{H}_{\text{cost}}$, most of which are degenerate. Eigenstates satisfying all constraints are shown in green, and constraint-violating states in red. Solid black lines mark the onset of a degenerate, constraint-violating ground state, where the minimum gap vanishes, here at $\lambda = 0.25$.

Since $\hat{H}_{\text{cost}}$ is diagonal in the computational basis, its eigenstates are basis states. On a constraint-satisfying basis state $|s\rangle$, the diagonal penalty contribution vanishes by construction, so that $\langle s|\hat{H}_{\text{cost}}|s\rangle = \langle s|\hat{H}_{\text{problem}}|s\rangle$ is independent of λ . On a constraint-violating basis state, by contrast, $\langle s|\hat{H}_{\text{constraints}}|s\rangle > 0$, and the energy grows linearly with λ . Both observations underlie the two regimes visible in Fig. 5.1.

For λ above the optimum, the low-energy spectrum is composed entirely of constraint-satisfying states, whose energies in physical units are independent of λ for the reason just stated. The maximum coupling, however, is dominated by the constraint contributions and therefore grows linearly in λ , $J \propto \lambda$, so that the dimensionless gap reduces as $\Delta E/J \propto \lambda^{-1}$. Choosing the penalty unnecessarily strong therefore wastes coupling 'budget' and compresses

the dimensionless spectrum that controls the annealing performance.

For λ below the optimum, conversely, constraint-violating eigenstates descend through the low-energy spectrum and progressively interleave with the constraint-satisfying ones. As long as the lowest constraint-violating state lies above the lowest constraint-satisfying excited state, the dimensionless gap continues to grow with decreasing λ , since the dilution of the encoding terms by the normalization is reduced. The largest gap is reached close to the value of λ where the lowest constraint-violating state and the lowest constraint-satisfying excited state become degenerate, in the example of Fig. 5.1b near $\lambda \approx 0.5$. Reducing λ further, as soon as a constraint-violating state crosses the ground state, the latter ceases to encode a feasible solution, and the minimum gap goes to zero; in the same example, this happens at $\lambda = 0.25$. The optimum value of λ is therefore set by the competition between the dilution of the encoding terms by the normalization at large λ and the loss of a constraint-satisfying ground state at small λ .

For other 3-SAT instances we observe a qualitatively similar behavior, with the threshold for the loss of feasibility occasionally located at higher values such as $\lambda \approx 0.5$, in which case the optimum lies even closer to the default choice $\lambda = 1$. Finding the optimal λ for a given instance generally requires knowledge of the optimal solution itself [357], so that λ must in practice be chosen by a problem-agnostic heuristic. The simple choice adopted here is sufficiently close to the optimum to ensure that our scaling analysis of the minimum gap (Sec. 5.1.4) is not significantly biased by suboptimal penalty strengths. More sophisticated approximations such as the one developed in Ref. [309] may yield additional improvements, but are unlikely to alter the qualitative conclusions of the PUBO-versus-QUBO comparison developed in the remainder of this section.

5.1.3 Structure of PUBO and QUBO formulations of 3-SAT

Before turning to the system-size scaling of the minimum gap in Sec. 5.1.4, we examine the structure of the low-lying instantaneous spectrum along the annealing path for typical instances generated by the three generators (see Sec. 3.2.5). The aim of this subsection is to give a qualitative picture of how the specific generator and the choice of PUBO versus QUBO formulation shape the gap profile.

Figure 5.2 shows the energies of the instantaneous ground state and first excited state along the standard annealing sweep, for one representative instance per generator at $N = 6$ variables, in both the native PUBO formulation and, where applicable, the corresponding QUBO reduction. Three qualitative features can be read off directly. First, both the **toughSAT** instance [Fig. 5.2a] and the **uniquePT1** instance [Fig. 5.2b] develop a pronounced gap minimum in the second half of the sweep. Second, the **uniquePT4** instance [Fig. 5.2c] exhibits no gap closing whatsoever and its first excited state remains separated from the ground state by an energy of order J throughout the sweep, in line with the discussion in Sec. 5.1.1: all three-body and two-body contributions cancel by construction and the cost Hamiltonian reduces to a sum of independent local fields with a ground state which can be read off trivially. Third, comparing the upper and lower panels for **toughSAT** and **uniquePT1**, the gap minima of the QUBO reductions are visibly smaller and located at later times along the sweep as in the corresponding PUBO formulations, illustrating the qualitative effect of the reduction at the level of a single instance.

The trivial nature of **uniquePT4** instances, confirmed by the gap profile in Fig. 5.2c, is a generic feature rather than an artifact of the specific instance shown. We therefore exclude this generator from the quantitative analysis below and restrict our PUBO-versus-QUBO comparison to **toughSAT** and **uniquePT1**.




Figure 5.2: Typical gap profiles of PUBO and QUBO formulations of randomly generated 3-SAT instances. The plots show the energies of the ground state (blue) and first excited state (orange) as a function of time during a standard annealing sweep according to Eq. (3.8) with driving strength h_x equal to the characteristic energy scale J of the cost Hamiltonian. The upper panels (a) and (b) depict the native PUBO formulations (cost Hamiltonian with up to three-body interactions, $J = \max_{ijk} |J_{ijk}^{(3)}|$), while the lower panels (d) and (e) show the associated QUBO reductions (cost Hamiltonian with at most two-body interactions, $J = \max_{ij} |J_{ij}^{(2)}|$). Note that two-body interactions are typically much stronger than three-body interactions. A direct decomposition of a 3-qubit unitary gate gives a ratio of $J^{(3)}/J^{(2)} \approx 1/4$ (see Sec. 5.1.5). The 3-SAT instances comprise $N = 6$ classical variables and are generated by the **toughSAT** generator with $M = \text{round}(4.24N) = 25$ clauses [panels (a) and (d)], the **uniquePT1** generator with $M = N + 6$ clauses [panels (b) and (e)], and the **uniquePT4** generator with $M/N = 4$ clauses [panel (c)]. The QUBO reductions of the instances in (a) and (b) require four and three additional slack variables (ancillary qubits), respectively. For instances generated by **uniquePT4** (c), no gap closing occurs since the generated instances are trivial (all three-body and two-body interactions cancel, $J = \max_i |h_i^z|$). The table (f) shows the mean minimum energy gap and its standard error of the mean for an ensemble of 200 instances corresponding to the scenarios in (a), (b) and (d), (e).

To make the comparison quantitative at fixed system size, we average the minimum gap along the annealing sweep over an ensemble of 200 random instances per generator and per formulation. The resulting mean minimum gaps and the error of the mean are reported in Fig. 5.2f. At $N = 6$, both generators yield mean gaps of comparable magnitude in the PUBO formulation, with **uniquePT1** slightly larger than **toughSAT**. This near-equality at small N is at first sight counterintuitive, since **uniquePT1** produces only $M = N + 6$ clauses against the $M \approx 4.24N$ clauses of **toughSAT**, and the QUBO reduction therefore requires fewer ancillary qubits for **uniquePT1** (three versus four for the instances shown in Fig. 5.2). The picture changes at larger system sizes, where the gap of **uniquePT1** instances closes significantly more rapidly than that of **toughSAT** instances, despite the smaller ancilla overhead, as we show in Sec. 5.1.4.

A feature worth noting is that the relative statistical uncertainty of the mean QUBO gap is noticeably larger than that of the corresponding PUBO gap. The reason is that, for a given

N , the number of ancillary qubits required by the brute-force quadratization fluctuates from instance to instance, depending on how many distinct cubic terms a given instance contains and how often they share pairs of variables. This instance-to-instance variability of the QUBO problem size translates directly into additional spread in the minimum gap and is absent in the PUBO formulation, where every instance with N variables uses exactly N qubits.




Figure 5.3: Distribution of the minimum energy gap and the number of third-order terms in the PUBO cost function (nonzero three-body interaction coefficients in the cost Hamiltonian). (a) Two-dimensional histogram, computed for 10 000 random 3-SAT instances generated via `toughSAT` with $N = 6$ and $M = 25$. (b)–(c) Marginal histograms of the minimum energy gap and the number of third-order terms, respectively, with fitted Gaussian distributions. No significant correlations can be discerned.

Correlation between gap and number of three-body terms. A natural question motivated by the per-instance variability of the QUBO reduction is whether, already in the native PUBO formulation, the minimum gap can be predicted from simple structural properties of the instance. Among such properties, an immediate one is the number of distinct three-body interaction coefficients in the cost Hamiltonian, that is, the number of cubic terms appearing in the underlying 3-SAT cost function. We address this question in Fig. 5.3, which shows the joint distribution of the minimum gap and the number of three-body terms over 10 000 random instances generated by `toughSAT` at $N = 6$ and $M = 25$. Both marginals are broadly distributed and approximately bell-shaped, with the number of three-body terms spread in the range 4 – 19. The two-dimensional histogram displays no discernible correlation between the two quantities. We conclude that the minimum gap of a




Figure 5.4: Scaling of mean minimum energy gap $\overline{\Delta E}$ with the number of Boolean variables N for ensembles of 200 randomly generated 3-SAT instances. The error bars show the statistical standard error of the mean, which is smaller than the marker size for most data points. The dashed lines show exponential fits to the data (ignoring $N = 3$) according to $\overline{\Delta E} = \epsilon e^{-\alpha N}$ with fit parameters ϵ and α given in Tab. 5.1. For problems generated with uniquePT1 (a), the gap closes approximately at the same rate in both the PUBO and QUBO formulation, while for toughSAT (b), the QUBO formulation exhibits a faster gap closing with increasing N . In all cases, the minimum gap relative to the characteristic energy scales is larger for the PUBO formulation ($J \sim J^{(3)}$) than for the corresponding QUBO reduction ($J \sim J^{(2)}$).

typical toughSAT instance cannot be inferred from a simple count of cubic terms in the cost function.

5.1.4 Scaling of the minimum energy gap

Table 5.1: Parameters of the fits shown in Fig. 5.4, characterizing the exponential scaling of the mean minimum energy gap $\overline{\Delta E} = \epsilon e^{-\alpha N}$ for randomly generated 3-SAT instances. The uncertainty of the fit parameters corresponds to the standard error of the mean for 200 realizations.

generator	PUBO		QUBO	
	ϵ_P/J	α_P	ϵ_Q/J	α_Q
toughSAT	0.306(8)	0.086(4)	0.17(2)	0.33(2)
uniquePT1	1.72(1)	0.352(1)	0.24(1)	0.364(8)

As established in Sec. 3.1.3, the minimum spectral gap encountered along the annealing sweep controls the adiabaticity time through $T \propto \hbar V / \Delta E^2$, and for hard combinatorial

problems the gap is expected to close exponentially with the problem size, $\overline{\Delta E} = \epsilon e^{-\alpha N}$, as the signature of a first-order quantum phase transition. The adiabaticity time then grows as $T \propto e^{2\alpha N}$ up to polynomial corrections, so that the exponent α is the central figure of merit for the efficiency of the annealing protocol.

Figure 5.4 shows the mean minimum energy gap as a function of N for 3-SAT instances generated by **toughSAT** and **uniquePT1**, in both the native PUBO formulation and the equivalent QUBO reduction as described in Sec. 5.1.1. Each data point is averaged over 200 random realizations, and the dashed lines show fits to $\overline{\Delta E} = \epsilon e^{-\alpha N}$ ignoring the smallest size $N = 3$. The resulting fit parameters are listed in Table 5.1.

In all four cases, the data are consistent with an exponentially closing gap over the investigated range of system sizes, in line with 3-SAT being NP-complete. Two features stand out from the comparison.

First, the two generators yield clearly distinct gap-scaling exponents already in the PUBO formulation. The exponent of **toughSAT** is roughly four times smaller than that of **uniquePT1**, $\alpha_P \approx 0.09$ versus $\alpha_P \approx 0.35$, so that on average **toughSAT** produces easier, though still exponentially hard, instances. This refines the gap-profile comparison at $N = 6$ in Fig. 5.2, where the two generators yielded comparable absolute gaps: the scaling analysis shows that this near-equality at small sizes does not extend to large N .

Second, and more relevant for the PUBO-versus-QUBO comparison, the two generators differ in how the QUBO reduction modifies the gap-scaling exponent. For **uniquePT1** [Fig. 5.4(a)], the PUBO and QUBO exponents are approximately identical, so that the reduction acts predominantly through a suppression of the prefactor in units of the characteristic energy scale. For **toughSAT** [Fig. 5.4(b)], by contrast, the QUBO exponent is markedly larger than its PUBO counterpart, $\alpha_Q \approx 0.33$ versus $\alpha_P \approx 0.09$, so that the dimensionless gap closes nearly four times faster in N in the QUBO reduction. Since $\alpha_P < \alpha_Q$, this exponential separation will, for sufficiently large N , dominate any prefactor and render the PUBO formulation asymptotically advantageous regardless of the relative physical interaction rates.

A quantitative comparison of the actual time-to-solution requires accounting for the different characteristic energy scales of the PUBO and QUBO Hamiltonians: as discussed in Sec. 3.1.4, the dimensionful adiabaticity time $T_{\text{phys}} = \hbar T/J$ is bottlenecked by the maximum three-body coupling for the PUBO formulation and by the maximum two-body coupling for the QUBO reduction, with the former typically smaller than the latter on physical hardware. We defer this analysis to Sec. 5.1.6, after the implementation of higher-order interactions on quantum hardware has been reviewed in Sec. 5.1.5.

Within the dimensionless adiabaticity time $JT/\hbar$, however, the gap scaling already permits an order-of-magnitude comparison with classical algorithms. Assuming that the matrix element V in the adiabatic condition of Eq. (3.11) scales at most polynomially in N , the dimensionless adiabaticity time grows as $JT/\hbar \propto e^{2\alpha N}$. For **toughSAT** in PUBO form, this yields $JT/\hbar \propto 1.19^N$, to be compared with the best known classical algorithm for the Unique 3-SAT problem (the promise version of 3-SAT in which each instance has at most one solution), which runs in time $\propto 1.307^N$ [358]. For the harder **uniquePT1** instances, the PUBO scaling is $JT/\hbar \propto 2.02^N$, close to the 2^N of classical brute-force enumeration. While the adiabaticity time is only an order-of-magnitude estimate of the true quantum-annealing runtime and does not capture the full computational complexity, these scalings indicate that, at least for some families of 3-SAT instances, quantum annealing in the native PUBO formulation may be competitive with state-of-the-art classical algorithms.

5.1.5 Implementing PUBO Hamiltonians on quantum hardware




Figure 5.5: (a) The three-qubit gate $R_{ZZZ}(\theta) = \exp(-i4\theta\hat{s}_1^z\hat{s}_2^z\hat{s}_3^z)$ can be decomposed into four CNOT gates plus a single-qubit rotation. (b) Each CNOT gate can be further decomposed into one R_{ZZ} gate of fixed angle and multiple single-qubit rotations. The quantum gates in the decompositions are defined as $R_\alpha(\theta) = \exp(-i\theta\hat{s}^\alpha)$, with $\alpha \in \{X, Y, Z\}$, $R_{ZZ}(\theta) = \exp(-2i\theta\hat{s}^z\hat{s}^z)$, and H denoting the Hadamard gate.

In this section, we discuss possibilities to implement three-body interactions in order to realize third-order PUBO formulations in quantum hardware. The exploration of such multi-body interaction terms and their role in enhancing quantum annealing protocols has become an active area of research. As we will see in the next section, the constant overhead of directly implementing PUBO problems is more than compensated, already for small problem sizes, by the advantageous scaling of the minimal gap investigated in the previous section.

Several proposals exist to directly engineer native three-body spin interactions of the type $\hat{s}_i^z\hat{s}_j^z\hat{s}_k^z$ (equivalently $\hat{\sigma}_i^z\hat{\sigma}_j^z\hat{\sigma}_k^z$) on the hardware level. For trapped-ion architectures, a detailed comparison of such schemes, including their typical achievable interaction strengths, was provided in Sec. 2.3; the discussed approaches range from concurrent off-resonant driving of first and second sideband and pulse sequences based on spin-dependent displacements and squeezings of a phonon bus, to Floquet engineering and perturbative constructions from two-body terms. Beyond trapped ions, proposals for effective multi-body interactions on superconducting qubit architectures have also been put forward [194]. As discussed in Sec. 2.3, most available analog schemes share the common drawbacks that their implementation is somewhat complex and that the achievable higher-order interaction rates are considerably slower compared to those of two-body interactions. Moreover, the tunability of the relative strengths of three-body interactions between different triples of spins depends on the utilized scheme as well as hardware characteristics. Nonetheless, these techniques may be exploited in next-generation devices to realize analog quantum annealing with higher-order interactions that benefits from the advantages of direct PUBO encodings in terms of spatial and temporal resources.

While native higher-order interactions can also act as physical implementations of multi-qubit gates in digital quantum computers, due to the disadvantages mentioned above, most available platforms implement only single- and two-qubit gates natively. These can still realize higher-order interaction gates, required for PUBO formulations, with only polynomial overhead. An example is given in Fig. 5.5: a three-qubit gate of the form $R_{ZZZ}(\theta) = \exp(-i4\theta\hat{s}_1^z\hat{s}_2^z\hat{s}_3^z)$ can be synthesized from four CNOT gates, plus a single-qubit rotation

around the z -axis [see Fig. 5.5a]. Each CNOT gate can be further decomposed into

$$\begin{aligned} \text{CNOT}_{1,2} &= e^{-i\pi/4} \\ &\times R_{y_1}\left(-\frac{\pi}{2}\right) R_{x_1}\left(-\frac{\pi}{2}\right) R_{x_2}\left(-\frac{\pi}{2}\right) H_1 H_2 \\ &\times R_{zz}\left(\frac{\pi}{2}\right) H_1 H_2 R_{y_1}\left(\frac{\pi}{2}\right), \end{aligned} \quad (5.8)$$

where we used a two-qubit controlled rotation gate R_{zz} (native to the MAGIC platform, see Sec. 1.3.3), Hadamard gates H and single-qubit rotations around x and y -axis. We show the decomposition in Fig. 5.5b. The overall circuit for the three-qubit gate then uses four two-body gates of type R_{zz} with fixed angle $\pi/2$ and 29 single-qubit rotations. In this decomposition, two adjacent single-qubit gates could be further combined into only one single-qubit rotation. Similar decompositions are possible with other two-qubit gates like the Mølmer-Sørensen (MS) type gate R_{xx} (see Sec. 1.3.2). Note that in digital quantum annealing protocols, the three-body interaction strength in each Trotter step is controlled by the rotation angle θ . Since θ enters in Fig. 5.5a only via a single $R_Z(\theta)$ gate, the overall three-qubit gate time is practically the same for different rotation angles.

As these examples show, the special structure of the ZZZ gate enables implementations with far fewer resources than necessary for a general three-body unitary operation (for example, state-of-the-art numerical sequential optimization requires 15 CNOT gates to decompose a general three-body unitary operation [359, 360]). Importantly, while such an overhead may still be non-negligible in current noisy intermediate-scale quantum (NISQ) devices, it is a rather small constant factor. As discussed below, this overhead can be more than compensated by the higher resource efficiency of the PUBO encoding.

5.1.6 Speedup of PUBO versus QUBO formulation

We now combine the dimensionless gap-scaling analysis of Sec. 5.1.4 with the hardware overhead estimated in Sec. 5.1.5 to compare the time-to-solution of the two formulations. Starting from the adiabaticity time $T = \hbar V / \Delta E^2$ of Eq. (3.12) and using the exponential gap closure $\overline{\Delta E} = \epsilon e^{-\alpha N}$, the dimensionful adiabaticity times of the PUBO and QUBO formulations read

$$T_P = \frac{\hbar}{J^{(3)}} \frac{\tilde{V}_P}{\tilde{\epsilon}_P^2} e^{2\alpha_P N}, \quad T_Q = \frac{\hbar}{J^{(2)}} \frac{\tilde{V}_Q}{\tilde{\epsilon}_Q^2} e^{2\alpha_Q N}, \quad (5.9)$$

with the dimensionless quantities $\tilde{\epsilon}_{P,Q} = \epsilon_{P,Q} / J^{(3,2)}$ and $\tilde{V}_{P,Q} = V_{P,Q} / J^{(3,2)}$ measured in units of the relevant characteristic energy scale of each formulation (see Sec. 3.1.4). Their ratio,

$$\frac{T_Q}{T_P} = \frac{J^{(3)}}{J^{(2)}} \frac{\tilde{V}_Q}{\tilde{V}_P} \left(\frac{\tilde{\epsilon}_P}{\tilde{\epsilon}_Q} \right)^2 e^{2(\alpha_Q - \alpha_P)N}, \quad (5.10)$$

is the central figure of merit: a value $T_Q/T_P > 1$ implies a faster anneal in PUBO. The four factors capture, respectively, the hardware overhead of three-body over two-body interactions (which disfavors PUBO), a generally polynomial dependence of the matrix-element ratio on N , the relative size of the dimensionless gap (typically larger for PUBO due to the absence of the energy-scale dilution induced by the QUBO penalty terms, see Sec. 5.1.2), and the relative gap-scaling exponent, which dominates the comparison for large system sizes whenever $\alpha_P \neq \alpha_Q$.

For the construction of three-body interaction gates discussed in the previous section, the synthesis of an $R_{ZZZ}(\theta)$ gate requires four two-qubit gates, so that, neglecting single-qubit gate times, $J^{(3)}/J^{(2)} \approx 1/4$. The matrix-element ratio is determined numerically; for the investigated system sizes we find $\tilde{V}_Q/\tilde{V}_P \approx 0.5$, with only a weak dependence on N . Inserting these values together with the gap-scaling parameters of Tab. 5.1 into Eq. (5.10) gives, for **uniquePT1**, $T_Q/T_P \approx 6.4 e^{0.024N}$. Since here $\alpha_P \approx \alpha_Q$, the exponential factor is essentially vanishing over the investigated range and the speedup is set by the prefactor: the ratio $(\tilde{\epsilon}_P/\tilde{\epsilon}_Q)^2 \approx 51$ more than compensates the fourfold hardware overhead, yielding a roughly constant PUBO speedup of about a factor of six. For **toughSAT**, the well separated exponents $\alpha_P \approx 0.086$ and $\alpha_Q \approx 0.33$ turn the comparison genuinely exponential, $T_Q/T_P \approx 0.4 e^{0.49N}$. Already at $N = 11$ classical variables, the speedup reaches roughly two orders of magnitude, growing by another order of magnitude for every ≈ 4.7 additional classical variables in the problem.

The estimates above depend on the choices of driving strength h^x and penalty strength λ in the QUBO reduction. The analyses of Secs. 3.1.5 and 5.1.2 place our default values, $h^x = J$ and $\lambda = 1$, close to the optimum, so that more sophisticated penalty heuristics such as the one of Ref. [309] are expected to modify the result only mildly. A polynomial scaling of $\tilde{V}_Q/\tilde{V}_P$ with N , which we cannot exclude beyond the investigated sizes, would moreover be sub-leading to the exponential factor in the **toughSAT** case and would not alter the qualitative conclusion. Within these caveats, the direct PUBO implementation yields a substantial, and for some instance families exponential, reduction of the time-to-solution compared to the conventional QUBO reduction.

Beyond the specific case of 3-SAT, similar advantages may occur for the broader class of problems with a natural polynomial formulation, including the examples collected in Sec. 3.2.6, all of which will incur a substantial qubit overhead under QUBO reduction. The case for direct PUBO implementations relies on continued progress in engineering native higher-order interactions, or in the case of digital implementations, supporting larger circuit depths for the respective gate decompositions.

5.2 Inhomogeneous driving

Having focused so far on modifications of the cost Hamiltonian, we now give a brief overview of the second avenue for increasing annealing performance mentioned in the introduction of this chapter, which acts instead on the driving Hamiltonian. In the standard annealing protocol (Sec. 3.1.2), the transverse-field driver is ramped down uniformly for all qubits. In inhomogeneous driving, or inhomogeneous quantum annealing (IQA), the schedule is modulated by qubit-dependent amplitudes,

$$\hat{H}_{\text{drive}}(t) = - \sum_i A_i(t) \hat{\sigma}_i^x, \quad (5.11)$$

with $A_i(0) = h^x$ and $A_i(t_{\text{final}}) = 0$ for all i . Letting different qubits cross the critical region (i.e. phase transition) at different times can enlarge the minimum spectral gap along the annealing evolution, which is a critical bottleneck for annealing protocols. All the schemes discussed below fit the form of Eq. (5.11), but target qualitatively different physical obstructions through different choices of $A_i(t)$.

Ref. [361] proposed driving a weakly disordered one-dimensional transverse-field Ising chain with a smooth spatial front that sweeps along the chain at velocity v and with characteristic width α . Because only a bounded number of spins pass through the critical region at

any given time, the residual energy after the sweep is suppressed by several orders of magnitude relative to the standard homogeneous ramp. A complementary analysis in Ref. [362] generalizes the construction to strongly disordered chains.

A second scheme, introduced in Ref. [285], specifically targets the first-order quantum phase transition that plagues standard annealing of the p -body mean-field ferromagnetic Ising model (see also discussion in Sec. 3.2.1). The transverse field is turned off site by site as the annealing progresses, which formally corresponds to the sharp-front limit ($\alpha \rightarrow 0$) of Ref. [361]. In the resulting family of annealing paths, the line of first-order transitions can be entirely avoided, implying an exponential speedup in the idealized mean-field limit. The same conclusion holds for the random-field variant, for which even non-stoquastic drivers (see Sec. 5.3) fail to remove the transition. Under finite temperature or a non-strictly sequential turn-off, however, new first-order transitions reappear and only quantitative improvements in the form of narrower free-energy barriers survive [363].

A third, experimentally oriented approach was demonstrated in Ref. [286] on the D-Wave 2000Q device. Each qubit is assigned an anneal offset δ_i that shifts its schedule along the global control axis of the driver, so that both the transverse-field amplitude and the problem-Hamiltonian coefficients at qubit i become functions of $s + \delta_i$ rather than s . The offsets are chosen according to the average effective field on qubit i ,

$$|\overline{\mathcal{F}}_i| = \frac{1}{2^{N_i}} \sum_{s_{j_1}, \dots, s_{j_{N_i}} \in \{\pm 1\}} \left| h_i + \sum_{j \in \mathcal{N}(i)} J_{ij} s_j \right|, \quad (5.12)$$

where $\mathcal{N}(i)$ is the set of N_i qubits coupled to i . Qubits with a large effective field, which belong to strongly coupled clusters and tend to freeze into a classical configuration before the rest of the system has converged, are assigned a delayed schedule. On Weak-Strong-Cluster and spin-glass benchmarks, this heuristic was reported to enhance the ground-state success probability by up to several orders of magnitude compared with the homogeneous protocol. Unlike the two schemes above, the inhomogeneity here is set by the problem Hamiltonian itself rather than by an external spatial coordinate.

Two caveats temper these results. First, a recent dynamical analysis shows that sequential turn-off IQA is considerably slower than the phase diagram suggests [364]: in mean-field models, the time needed to reach a fixed residual-energy threshold exceeds the adiabatic estimate by a factor of N , and in generic non-mean-field models first-order transitions with an exponentially vanishing gap are argued to be common, which prevents IQA from reaching low-energy states on any timescale. Second, the effective-field heuristic in Eq. (5.12) involves a sum over 2^{N_i} neighbor configurations and becomes prohibitive for densely connected instances.

5.3 Non-stoquastic quantum annealing

After modifying the cost Hamiltonian in Sec. 5.1 and the driving Hamiltonian in Sec. 5.2, the third avenue discussed in this chapter consists of augmenting the annealing Hamiltonian with a new term, active only during the intermediate stages of the sweep. The standard annealing Hamiltonian of Eq. (3.8) has a special structural property in the computational basis: all its off-diagonal matrix elements are real and non-positive. This feature turns out to have implications both for the classical simulability of the annealing dynamics and for the range of computational problems that can plausibly benefit from quantum speedups. In this section we review the corresponding concept of stoquasticity and then describe how

the standard protocol can be enriched by a non-stoquastic *catalyst* term. Specific models in which such catalysts play a central role will be studied in Chapter 6.

Stoquastic Hamiltonians

A Hamiltonian is called *stoquastic* in a given local basis if all its off-diagonal matrix elements in that basis are real and non-positive [190]. By the Perron-Frobenius theorem, the ground state of such a Hamiltonian can then be chosen with non-negative amplitudes in the same basis, and the associated partition function admits a decomposition into non-negative weights. The standard transverse-field Ising Hamiltonian of Eq. 3.8 is stoquastic in the computational basis: the cost Hamiltonian is diagonal, while the transverse-field driver contributes only negative off-diagonal entries $-h^x$.

The relevance of stoquasticity stems from its connection to the sign problem in quantum Monte Carlo methods (QMC) [365]. For a stoquastic Hamiltonian, the Suzuki-Trotter decomposition of the partition function yields a classical statistical-mechanics problem with non-negative Boltzmann weights, which can in principle be sampled by standard Monte Carlo methods. Non-stoquastic Hamiltonians, in contrast, generate weights of oscillating sign whose cancellations can cause the statistical error to grow exponentially with the system size, rendering generic QMC algorithms inefficient.

Stoquasticity is a basis-dependent property, and in some cases a Hamiltonian that is non-stoquastic in one basis can be rendered stoquastic by a local unitary transformation. Determining whether such a sign-curing transformation exists is itself computationally hard in general [366, 367], which limits its utility as a practical diagnostic.

The absence of a sign problem does not imply efficient classical simulation: the ground-state energy problem for classical Ising models, which are trivially stoquastic (the Hamiltonian only has diagonal entries), is already NP-hard. In the framework of Hamiltonian complexity [36, 37], the k -local Hamiltonian problem restricted to stoquastic instances is contained in the probabilistic class AM [190], and complete for the intermediate class StoqMA [368]. In the frustration-free case, stoquastic LH-MIN is MA-complete [369]. These containments place stoquastic LH-MIN below the QMA-complete general local Hamiltonian problem under standard complexity-theoretic assumptions, and provide the complexity-theoretic justification for the expectation that non-stoquastic Hamiltonians unlock computational resources inaccessible to the standard protocol. Nonetheless, recent work has shown that this expectation should not be pushed too far: Refs. [370, 371] demonstrated (sub)exponential oracle separations between adiabatic computation with no sign problem and classical computation, so that stoquastic adiabatic computation is, by itself, already strictly more powerful than classical computation in the oracle model. A concrete result in favor of non-stoquasticity is given in Ref. [191]: the transverse-field Ising model supplemented by two-body $\hat{\sigma}_i^x \hat{\sigma}_j^x$ interactions, as well as a second model consisting of longitudinal and transverse fields with two-body $\hat{\sigma}_i^z \hat{\sigma}_j^x$ couplings, are among the simplest known 2-local qubit Hamiltonians whose local Hamiltonian problem is QMA-complete. Combined with the equivalence between adiabatic and circuit-model quantum computation established in Ref. [32], these Hamiltonians are therefore universal for adiabatic quantum computation.

Catalyst Hamiltonians

A natural strategy to exploit the additional computational resources associated with non-stoquasticity, without giving up the easily prepared initial state of the standard protocol, is to introduce a *catalyst* Hamiltonian $\hat{H}_{\text{catalyst}}$ that is active only during the intermediate

stages of the anneal [30, 183]. The modified annealing Hamiltonian reads

$$\tilde{H}(s) = (1 - s) \hat{H}_{\text{drive}} + \lambda(s) \hat{H}_{\text{catalyst}} + s \hat{H}_{\text{cost}}, \quad (5.13)$$

where the schedule $\lambda(s)$ satisfies $\lambda(0) = \lambda(1) = 0$. With this choice, both the initial and final states of the annealing protocol remain those of the original stoquastic Hamiltonian, and the non-stoquastic contribution is confined to an intermediate window where the gap-closing bottleneck typically occurs.

The common choice of catalyst are antiferromagnetic two-body XX interactions,

$$\hat{H}_{\text{catalyst}} = + \sum_{i < j} K_{ij} \hat{\sigma}_i^x \hat{\sigma}_j^x, \quad (5.14)$$

with $K_{ij} > 0$, which renders the total Hamiltonian non-stoquastic in the computational basis. More generally, higher-order transverse terms such as $\hat{\sigma}_i^x \hat{\sigma}_j^x \hat{\sigma}_k^x$ can also be considered [372]. The physical effect of such a catalyst is twofold. First, by introducing quantum transitions that flip multiple spins simultaneously, it modifies the spectrum along the annealing path and, in favorable cases, enlarges the minimum gap, thus relaxing the adiabatic condition [195, 373]. Second, even when the minimum gap is not significantly increased, the catalyst can reshape the instantaneous spectrum to create controlled avoided crossings that can be traversed *diabatically* rather than adiabatically, exploiting rather than avoiding non-adiabatic transitions [374–376].

The paradigmatic success of the catalyst approach is the infinite-range ferromagnetic p -spin model, where the standard annealing protocol is obstructed by a first-order quantum phase transition with an exponentially closing gap. Seki and Nishimori [195] showed analytically that adding an antiferromagnetic XX catalyst with a suitable amplitude converts the first-order transition into a second-order one, replacing the exponentially closing gap with a polynomially closing gap and thus implying an exponential speedup. The result was subsequently extended to many-body transverse interactions [372] and to the Hopfield model with a finite number of patterns [373]. An overview is given in Ref. [196]. Building on this line of work, Ref. [294] constructed explicit 2-local and geometrically local Ising models that display exponential advantages from non-stoquastic catalysts even outside the mean-field limit. We will use this geometrically local Ising model, as well as the p -spin model, as benchmark problems in Chapter 6.

The picture is, however, not uniformly positive. Numerical studies of random Ising spin glasses in Ref. [197] found that non-stoquastic drivers outperform their stoquastic counterparts only on a small fraction of hard instances, and attributed the observed improvements to catalyst-induced diabatic transitions rather than to a systematic enlargement of the spectral gap. Ref. [377] subsequently proved that for adiabatic paths ending in a classical Hamiltonian, a generic non-stoquastic path has a *smaller* minimum gap than its stoquasticized counterpart obtained by flipping the signs of off-diagonal matrix elements. This establishes that non-stoquasticity is not a generic source of speedups, and that the known advantages, such as the results for the p -spin model, rely on special structural features of the problem Hamiltonian. Finally, Ref. [378] showed that first-order quantum phase transitions of topological origin, in a strongly frustrated Ising ladder, cannot be removed by either stoquastic or non-stoquastic XX catalysts, which indicates that the p -spin model mechanism does not carry over to all first-order transitions.

Despite this mixed picture, non-stoquastic catalysts remain one of the most actively investigated routes to improve the stoquastic standard annealing protocol, both because of the provably exponential speedups in specific models and because the required XX and

XZ interactions are experimentally accessible on current hardware platforms. The following chapter studies non-stoquastic annealing in detail for two benchmark models, focusing on the interplay between catalyst strength, the induced gap structure and quantum resources.

Chapter 6

Quantum resources in non-stoquastic annealing

Quantum annealing prepares the ground state of a problem Hamiltonian via adiabatic interpolation from a simple initial state, with a runtime controlled by the inverse square of the minimum spectral gap encountered along the evolution (see Sec. 3.1.3). In its standard transverse-field implementation, the annealing Hamiltonian is stoquastic, free of the sign problem [190, 365]. Arguments from complexity theory indicate that non-stoquasticity may be necessary to achieve a quantum advantage through adiabatic computation [37, 190, 369], and explicit constructions show that suitably designed non-stoquastic catalyst Hamiltonians (Sec. 5.3) can substantially modify the spectrum along the protocol. A paradigmatic example is the ferromagnetic p -spin model, for which an antiferromagnetic transverse catalyst converts the first-order quantum phase transition of the stoquastic protocol, with its exponentially closing gap, into a second-order one with polynomial gap scaling, yielding an exponential reduction of the adiabatic runtime [195, 196].

A larger spectral gap is necessary but not sufficient to guarantee a quantum advantage. The instantaneous ground state can also be tracked classically whenever it admits an efficient tensor-network description, or whenever it lies close to the set of stabilizer states accessible to stabilizer-tableau methods [103]. The corresponding obstructions are quantified by two complementary quantum resources of the state, the bipartite entanglement entropy and the nonstabilizerness, or magic. For stoquastic protocols, both have been shown to develop pronounced local maxima along the sweep, in the form of entanglement [379, 380] and magic [381] barriers. Whether non-stoquastic catalysts enhance or reduce these resources, and how their behavior correlates with the spectral improvement reported in Ref. [195], has so far remained largely unexplored.

In this chapter, we address this question by jointly monitoring, along the annealing evolution, the energy gap (and in particular its minimum along the sweep), the bipartite entanglement entropy, and the second stabilizer Rényi entropy [117] of the instantaneous ground state, for the two benchmark models introduced in Secs. 3.2.1 and 3.2.2: the fully connected ferromagnetic p -spin model, whose permutation symmetry enables efficient evaluation, and the geometrically local Ising model of Ref. [294], which is more suitable for near-term hardware implementations. After reviewing the two resources of interest and the numerical methods used to evaluate them efficiently (Sec. 6.1), we analyze the p -spin model in Sec. 6.2 and the local Ising model in Sec. 6.3, including a microscopic interpretation of the plateau structure induced by the catalyst in the latter case. In both models, we find

that the protocols with larger non-stoquasticity exhibit entanglement and stabilizer Rényi entropy that are both larger in magnitude and qualitatively different in behavior, with both resources displaying local peaks near the spectral bottleneck. The performance advantage of non-stoquastic annealing thus correlates with the production of states that resist both tensor-network and Clifford-based classical simulation.

The content of this chapter is based on Ref. [204].

6.1 Entanglement and non-stabilizerness

To characterize the quantum resources generated during non-stoquastic annealing protocols, we monitor two complementary resources, the bipartite entanglement entropy and the stabilizer Rényi entropy (SRE). The former quantifies the quantum correlations established between subsystems of a pure state, while the latter measures its non-stabilizerness, or magic, namely the distance of the state from the set of classically simulable stabilizer states. Both quantities are key ingredients of any quantum computational advantage, since they obstruct complementary families of efficient classical simulation methods. States with limited bipartite entanglement admit compact descriptions in terms of matrix product states (MPS) and more general tensor network methods, and can therefore be simulated efficiently on a classical computer. Stabilizer states, including those with volume-law entanglement, admit an efficient classical description; Clifford circuits acting on stabilizer states can therefore be simulated in polynomial time by the Gottesman–Knill theorem [103, 382]. Evading both classes of simulation strategies therefore requires the simultaneous generation of extensive entanglement and non-stabilizerness along the annealing trajectory.

6.1.1 Bipartite entanglement entropy

For a pure quantum state $|\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B$ of a bipartite system, the bipartite entanglement entropy is defined as the von Neumann entropy of the reduced density matrix $\rho_A = \text{Tr}_B |\psi\rangle\langle\psi|$,

$$S(\rho_A) = -\text{Tr}(\rho_A \log_2 \rho_A) = -\sum_i \lambda_i \log_2 \lambda_i, \quad (6.1)$$

where $\{\lambda_i\}$ are the eigenvalues of ρ_A . Equivalently, the λ_i are the squared Schmidt coefficients of $|\psi\rangle$ with respect to the bipartition $A|B$, so that $S(\rho_A) = S(\rho_B)$. For subsystems of dimensions d_A and d_B , the entropy is bounded by $\log_2 \min(d_A, d_B)$, a bound which is saturated on maximally entangled states. In the following we always adopt a balanced bipartition, $n_A = n_B = N/2$, for which $d_A = d_B = 2^{N/2}$ and the maximal entropy is $N/2$.

Numerical evaluation. For generic N -qubit states without exploitable symmetries, the Schmidt coefficients are most efficiently extracted by reshaping the 2^N amplitudes of $|\psi\rangle$ into a $2^{n_A} \times 2^{n_B}$ matrix Ψ and performing a singular value decomposition, $\Psi = U \Sigma V^\dagger$. The diagonal entries of Σ are the Schmidt coefficients and directly yield the eigenvalues λ_i entering Eq. (6.1). This procedure bypasses the explicit construction of ρ_A and its diagonalization, but is still limited by the cost of the SVD, which scales as $\mathcal{O}(2^{3N/2})$ for a balanced bipartition.

The computational effort is drastically reduced in the presence of permutation symmetry, as in the fully connected p -spin model studied in Sec. 3.2.1. For a Hamiltonian invariant under permutations of the qubits the ground state during an annealing protocol remains within the fully symmetric subspace of total spin $S = N/2$, which is spanned by the $N + 1$

Dicke states $|S, m\rangle$ with $m = -S, \dots, S$. Any state within this subspace can be expanded as

$$|\psi\rangle = \sum_{m=-S}^S c_m |S, m\rangle, \quad (6.2)$$

so that its complete specification requires only $N+1$ coefficients instead of 2^N . The Schmidt decomposition of each Dicke state with respect to a splitting into n_A and n_B qubits is known in closed form [383], and implies that the reduced density matrix can be written as $\rho_A = AA^\dagger$, with

$$A_{m,j} = c_{m+j} \sqrt{\frac{\binom{n_A}{m} \binom{n_B}{j}}{\binom{N}{m+j}}}, \quad m = 0, \dots, n_A, \quad j = 0, \dots, n_B. \quad (6.3)$$

Here m and j count the number of spin excitations in subsystems A and B , respectively, and $m+j$ is the total excitation number labeling the Dicke component of $|\psi\rangle$. Eq. (6.3) follows from the combinatorics of distributing $m+j$ excitations among the two halves. Diagonalizing the $(n_A+1) \times (n_A+1)$ matrix ρ_A then yields the full Schmidt spectrum at a cost that is only polynomial in N , making the evaluation of Eq. (6.1) tractable beyond the regime accessible to a brute-force SVD on the full 2^N -dimensional state vector.

6.1.2 Stabilizer Rényi entropy

Non-stabilizerlessness can be quantified through the stabilizer Rényi entropy (here of second order), introduced in Ref. [117]. For a pure N -qubit state $|\psi\rangle$ it is defined as

$$M_2(|\psi\rangle) = -\log_2 \left[\frac{1}{2^N} \sum_{P \in \mathcal{P}_N} \langle \psi | P | \psi \rangle^4 \right], \quad (6.4)$$

where the sum runs over all 4^N phase-free elements of the N -qubit Pauli group $\mathcal{P}_N = \{\sigma_{i_1} \otimes \dots \otimes \sigma_{i_N}\}$ with $\sigma_i \in \{I, X, Y, Z\}$. The rescaled distribution $\Xi_P(|\psi\rangle) = \langle \psi | P | \psi \rangle^2 / 2^N$ is non-negative and normalized to unity on pure states, and M_2 coincides with the classical Rényi-2 entropy of Ξ_P shifted by $\log_2 2^N$. The SRE is non-negative, invariant under Clifford unitaries, and vanishes if and only if $|\psi\rangle$ is a stabilizer state [117, 384]. This reflects the fact that stabilizer states have sharply peaked Pauli distributions, with $\langle \psi | P | \psi \rangle = \pm 1$ on the 2^N elements of the associated stabilizer group and zero elsewhere, whereas generic states spread their weight over exponentially many Pauli strings.

Numerical evaluation. The direct evaluation of Eq. (6.4) on a generic many-body state requires the computation of all 4^N Pauli expectation values. Each expectation value costs $\mathcal{O}(2^N)$ operations on a dense state vector, yielding an overall scaling of $\mathcal{O}(8^N)$. The SRE therefore becomes prohibitively expensive at considerably smaller system sizes than the entanglement entropy, whose brute-force evaluation scales as $\mathcal{O}(2^{3N/2})$.

The situation is again vastly improved for permutation-symmetric states. For such states the expectation value of a Pauli string depends only on the numbers of X , Y , Z , and identity factors it contains, and not on the specific qubits on which they act. The 4^N Pauli operators can therefore be grouped into equivalence classes labelled by quadruples $\mathbf{N} = (N_x, N_y, N_z, N_0)$ with $N_x + N_y + N_z + N_0 = N$, each class carrying a degeneracy

$g(\mathbf{N}) = N!/(N_x! N_y! N_z! N_0!)$. The number of inequivalent classes is

$$D = \binom{N+3}{3} = \mathcal{O}(N^3), \quad (6.5)$$

so that the sum in Eq. (6.4) can be evaluated in polynomial time by summing over equivalence classes weighted by their degeneracies, using the Dicke representation (6.2) of $|\psi\rangle$ [385]. This reduction makes the SRE accessible for larger system sizes, similar to the entanglement entropy, in permutation-symmetric models like the p -spin model (Sec. 3.2.1).

6.2 Numerical results for the p -spin model

In this section, we apply the non-stoquastic annealing protocol to the ferromagnetic p -spin model (introduced in Ref. [195]) and analyze the interplay between spectral properties and quantum resources during the protocol. After defining the catalyst Hamiltonian and the resulting two-parameter annealing schedule (Sec. 6.2.1), we examine the energy gap, the bipartite entanglement entropy, and the second stabilizer Rényi entropy across the annealing parameter space (Sec. 6.2.2) and study their finite-size scaling along sinusoidal schedules (Sec. 6.2.3). Finally, we analyze how entanglement and non-stabilizerness are distributed in the form of barriers along the annealing sweep (Sec. 6.2.4).

6.2.1 Non-stoquastic annealing protocol for the p -spin model

We recall from Sec. 3.2.1 the cost Hamiltonian of the ferromagnetic p -spin model,

$$H_{\text{cost}} = -N \left(\frac{1}{N} \sum_{i=1}^N \hat{\sigma}_i^z \right)^p, \quad (6.6)$$

together with the standard transverse-field driver $H_{\text{drive}} = -\sum_i \hat{\sigma}_i^x$. Under the conventional stoquastic schedule, the protocol encounters a first-order quantum phase transition, accompanied by an exponentially small minimum gap that imposes an exponential adiabatic runtime [192].

To circumvent this bottleneck, Seki and Nishimori [195] introduced antiferromagnetic transverse fluctuations through the non-stoquastic catalyst Hamiltonian

$$H_{\text{catalyst}} = +N \left(\frac{1}{N} \sum_{i=1}^N \hat{\sigma}_i^x \right)^2, \quad (6.7)$$

which corresponds to an all-to-all two-body $\hat{\sigma}_i^x \hat{\sigma}_j^x$ interaction. The full time-dependent annealing Hamiltonian then reads

$$\hat{H}(s, \lambda) = s\lambda H_{\text{cost}} + s(1-\lambda) H_{\text{catalyst}} + (1-s) H_{\text{drive}}, \quad (6.8)$$

with the two parameters $s, \lambda \in [0, 1]$ both varying in time. The protocol begins at $(s, \lambda) = (0, \text{arbitrary})$, where only the driver is active, and ends at $(s, \lambda) = (1, 1)$, where only the cost Hamiltonian remains. For $5 \leq p \leq 21$, paths through the (s, λ) plane exist that bypass all first-order transitions and instead cross a second-order transition with a polynomially closing gap, yielding an exponential improvement of the adiabatic runtime [195, 196]. In the limit $p \rightarrow \infty$, however, first-order transitions become unavoidable, consistent with the optimality of the Grover bound.

Because $\hat{H}_{\text{catalyst}}$ depends only on the collective transverse magnetization, it preserves the permutation symmetry already shared by $\hat{H}_{\text{cost}}$ and $\hat{H}_{\text{drive}}$. The full annealing Hamiltonian therefore commutes with all pairwise spin permutations at every (s, λ) , and the dynamics remains confined to the fully symmetric Dicke subspace of dimension $N + 1$. Following the construction detailed in Sec. 3.2.1, the cost and driving Hamiltonians in this subspace are given by Eqs. (3.23) and (3.24), while the catalyst, expressed in terms of the collective spin operator $\hat{S}^x$ of Eq. (3.21), reads

$$\hat{H}_{\text{catalyst}} = \frac{4}{N} (\hat{S}^x)^2, \quad (6.9)$$

where we used $\sum_i \hat{\sigma}_i^x = 2\hat{S}^x$. The full annealing Hamiltonian thus takes the compact form

$$\hat{H}(s, \lambda) = -\frac{2^p s \lambda}{N^{p-1}} (\hat{S}^z)^p + \frac{4 s (1 - \lambda)}{N} (\hat{S}^x)^2 - 2(1 - s) \hat{S}^x, \quad (6.10)$$

whose matrix elements in the Dicke basis follow directly from Eq. (3.25). Consequently, the numerical methods discussed in the previous section, namely exact diagonalization in the Dicke basis, evaluation of the bipartite entanglement entropy on a balanced bipartition $n_A = n_B = N/2$, and computation of the second stabilizer Rényi entropy via Pauli equivalence classes [385], apply directly. In the following we focus on $p = 5$, for which a non-stoquastic path with polynomial gap scaling is known to exist [195], and we access system sizes up to $N = 100$.

6.2.2 Quantum resources across the annealing parameter space

We start by analyzing the behavior of the three observables across the (s, λ) parameter space. The first row of Fig. 6.1 shows, for $N = 30$ and $p = 5$, the energy gap [panel (a)], the bipartite entanglement entropy [panel (b)], and the second stabilizer Rényi entropy [panel (c)], evaluated on the instantaneous ground state of $\hat{H}(s, \lambda)$ on a uniform 301×301 grid in $[0, 1] \times [0, 1]$. All three quantities are shown on a logarithmic scale to enhance the visibility of the relevant features.

The three observables share a striking qualitative similarity. The locus of the minimum energy gap (white dashed line in Fig. 6.1) coincides with the region where both the entanglement entropy and the SRE attain their maxima, indicating that quantum correlations and non-stabilizerness both peak in the critical region of the protocol. This pattern is consistent with the phase diagram of the p -spin model under non-stoquastic annealing (see Fig. 3 of Ref. [195]): the line of the minimum energy gap traces the boundary between the paramagnetic and the ferromagnetic phase in the (s, λ) plane. For large λ this boundary corresponds to a first-order transition, but as λ decreases it terminates at a critical point near $(s, \lambda) \approx (0.4, 0.3)$, beyond which the transition becomes second order and the exponentially small gap is avoided.

Away from the line of the minimum spectral gap, the entanglement entropy and the SRE display qualitatively distinct behaviors. For small s , both remain close to zero across all λ . After their respective maxima, the SRE decreases smoothly with s , whereas the entanglement entropy drops more sharply. In both cases, this variation becomes progressively smoother as the path moves toward the second-order region. While the two quantities thus respond to the same critical region, they encode complementary aspects of the underlying quantum structure of the ground state, namely the bipartite correlations between subsystems and the spread of the state over the Pauli basis [117].




Figure 6.1: **First row:** (a) energy gap, (b) bipartite entanglement entropy, and (c) second stabilizer Rényi entropy as functions of the annealing parameters s and λ , for the p -spin model with $N = 30$ and $p = 5$. All observables are evaluated on the instantaneous ground state of the Hamiltonian in Eq. (6.8) and shown on a logarithmic scale. **Second row:** finite-size scaling of (d) the minimum energy gap, (e) the maximum bipartite entanglement entropy, and (f) the maximum second stabilizer Rényi entropy along annealing paths defined by $\lambda(s) = 1 - A \sin(\pi s)$, for amplitudes $A = 0, 0.5, 0.7, 0.9$. For each protocol, the minimum gap and the maximum values of entanglement and SRE are extracted along the sweep and plotted as functions of system size N . Dashed lines in panel (d) correspond to fits of the minimum energy gap. For $A = 0$ and $A = 0.5$, exponential fits of the form $\Delta E_{\min} \approx \exp(a + bN)$ are used, with $(a, b)_{A=0} = (0.597, -0.134)$ and $(a, b)_{A=0.5} = (-0.666, -0.085)$. For $A = 0.7$ and $A = 0.9$, fits of the form $\Delta E_{\min} \approx \exp(a + bN + cN^2)$ are employed, with $(a, b, c)_{A=0.7} = (-1.34, -8.68 \times 10^{-2}, 5.43 \times 10^{-4})$ and $(a, b, c)_{A=0.9} = (-2.23, -9.46 \times 10^{-2}, 6.84 \times 10^{-4})$. Solid lines in panel (e) are guides to the eye. In panel (f), dashed lines show linear fits of the maximum SRE for $N \geq 30$ of the form $M_2^{\max} = a + bN$, with $(a, b)_{A=0} = (4.42, 0.033)$, $(a, b)_{A=0.5} = (3.87, 0.206)$, $(a, b)_{A=0.7} = (0.019, 0.415)$, and $(a, b)_{A=0.9} = (0.563, 0.409)$.

6.2.3 Finite-size scaling along sinusoidal annealing paths

Guided by the structure of the phase diagram discussed above, we now consider one-parameter families of annealing paths in which λ is varied as a function of s . Specifically, we adopt the sinusoidal schedule

$$\lambda(s) = 1 - A \sin(\pi s), \quad (6.11)$$

where the amplitude $A \in [0, 1]$ controls the strength of the non-stoquastic catalyst term along the protocol [cf. Eq. (6.8)]. This choice is representative of a broader class of admissible paths, the only requirements being that $\lambda(s = 0)$ is arbitrary and that $\lambda(s = 1) = 1$, so that the initial and final Hamiltonians coincide with H_{drive} and H_{cost} , respectively. Different

values of A correspond to qualitatively different paths in the phase diagram: $A = 0$ recovers the standard stoquastic schedule; $A = 0.5$ still crosses a first-order transition; $A = 0.7$ approaches the critical point at which the transition changes from first to second order; and $A = 0.9$ lies entirely in the second-order region, thereby avoiding all first-order transitions.

The second row of Fig. 6.1 [panels (d)–(f)] summarizes the finite-size behavior of the three observables along these four paths. From left to right, the panels show the minimum energy gap, the maximum bipartite entanglement entropy, and the maximum stabilizer Rényi entropy attained along the sweep, as functions of N . For the gap and the SRE we consider system sizes from $N = 10$ to $N = 60$ in steps of five, while for the entanglement entropy we extend the analysis up to $N = 100$ for the non-stoquastic paths $A = 0.7$ and $A = 0.9$ to better resolve their large- N behavior.

The annealing parameter s is discretized uniformly, with additional refinement in the vicinity of the relevant extrema, namely the minima of the energy gap and the maxima of the bipartite entanglement entropy and SRE. For the energy gap, we use $\Delta s = 10^{-7}$ at $A = 0$ and $\Delta s = 5 \times 10^{-6}$ at $A = 0.5, 0.7, 0.9$. For the bipartite entanglement entropy, we use $\Delta s = 5 \times 10^{-7}$ at $A = 0$ and $\Delta s = 5 \times 10^{-6}$ for the remaining amplitudes. For the largest sizes, $N = 70, 80, 90, 100$, this grid is further refined to $\Delta s = 10^{-7}$ at $A = 0.7$ and $A = 0.9$. For the SRE, whose evaluation is computationally more expensive, we use a slightly coarser grid, with $\Delta s = 5 \times 10^{-6}$ at $A = 0$ and $\Delta s = 5 \times 10^{-5}$ at the remaining amplitudes. We find that these discretizations give satisfactorily converged estimates of the extremal values of all observables considered.

The minimum gap exhibits two qualitatively distinct scaling behaviors, captured by the dashed fits in Fig. 6.1(d), which retain only the leading contributions on a logarithmic scale. For $A = 0$ and $A = 0.5$, corresponding to paths crossing a first-order transition, the gap decreases exponentially with system size, $\Delta E_{\min} \approx \exp(a + bN)$, in agreement with the standard expectation [192]. For $A = 0.7$ and $A = 0.9$, corresponding to paths that approach or lie within the second-order region, the decay is significantly slower and is well described on a logarithmic scale by a quadratic dependence, $\log \Delta E_{\min} \approx a + bN + cN^2$. This change of scaling reflects the softening of the spectral bottleneck once first-order transitions are avoided.

The maximum entanglement entropy [Fig. 6.1(e)] follows a different pattern. For $A = 0$ and $A = 0.5$, it grows rapidly with N at small system sizes and then saturates to a value close to unity (the small dip at the largest N is a finite-resolution artifact of the discretization in s). For $A = 0.7$, corresponding to a path close to the critical point between first- and second order transition, the maximum entanglement continues to increase over the entire range of system sizes investigated, suggesting stronger quantum correlations along the sweep. For $A = 0.9$, the growth is smoother and remains compatible with eventual convergence, although with a slower approach to the asymptotic value than in the first-order case.

The behavior of the SRE [Fig. 6.1(f)] is qualitatively different from both. For all four protocols, the maximum SRE grows approximately linearly with N , indicating an extensive scaling well described by $M_2^{\max} = a + bN$. The slope b , however, depends strongly on the path: paths crossing a first-order transition ($A = 0$ and $A = 0.5$) show a much weaker increase, whereas paths that approach or lie within the second-order region ($A = 0.7$ and $A = 0.9$) display a substantially steeper growth. Although SRE is therefore extensive for all protocols, it is strongly enhanced when first-order transitions are avoided. We further note that the slope is essentially identical between $A = 0.7$ and $A = 0.9$, suggesting that, once the path crosses the second-order region, the rate of non-stabilizerness generation no longer depends on the precise schedule.




Figure 6.2: **Top row:** bipartite entanglement entropy of the instantaneous ground state along the annealing sweep as a function of s , for sinusoidal paths $\lambda(s) = 1 - A \sin(\pi s)$ with $A = 0, 0.5, 0.7, 0.9$ (from left to right). For each amplitude, results are shown for system sizes $N = 20, 30, 40, 50$. **Bottom row:** stabilizer Rényi entropy density M_2/N for the same state along the same paths and system sizes. Insets show zoom-ins of the regions around the maximum entanglement and SRE values.

6.2.4 Entanglement and magic barriers

To resolve how entanglement and magic are distributed along the annealing sweep, we examine the corresponding profiles as functions of the annealing time s . Fig. 6.2 shows the bipartite entanglement entropy (top row) and the stabilizer Rényi entropy density M_2/N (bottom row) along the sinusoidal schedule of Eq. (6.11) for $A = 0, 0.5, 0.7, 0.9$ and $N = 20, 30, 40, 50$. The annealing time s is discretized as in the previous subsection, with finer resolution near the peaks and a coarser step $\Delta s = 0.002$ elsewhere.

The entanglement profile is dominated by pronounced peaks for $A = 0, 0.5$, and 0.7 , whose sharpness increases as A decreases and, at fixed A , as N grows. For $A = 0.9$ the structure changes qualitatively: the single peak splits into multiple local maxima distributed over a wider range of s , in line with the broader critical region traversed by paths that avoid first-order transitions. In all cases, the peak height of the entanglement entropy increases with system size.

A complementary trend is observed for the SRE density. As A increases, the profile becomes broader and smoother. For $A = 0$ and $A = 0.5$, the peak height of M_2/N decreases with N , indicating sub-extensive scaling of the total M_2 in this regime. For $A = 0.7$ and $A = 0.9$, by contrast, the peak height of the density remains approximately constant in N , consistent with the extensive scaling of the total SRE established in the previous subsection: a barrier of height $\sim \mathcal{O}(1)$ in M_2/N corresponds to a barrier of height $\sim \mathcal{O}(N)$ in M_2 .

These profiles motivate the term *barriers* for the peaks observed along the sweep. The instantaneous ground state must traverse a region of substantial entanglement and non-stabilizerness in order to interpolate between the trivial driver ground state at $s = 0$ and the polarized ground state of the cost Hamiltonian at $s = 1$, and the height and shape of these barriers reflect the type of transition crossed.

Taken together, the results of this section indicate that non-stoquastic driving reshapes

both the annealing performance and the structure of the generated quantum resources beyond what is captured by the energy gap alone. Along the protocols of interest, the bipartite entanglement entropy remains bounded for stoquastic protocols and grows larger and saturates more slowly for non-stoquastic protocols. The total SRE, on the other hand, grows extensively, with a linear slope that is largest for the protocols that avoid first-order transitions. This larger extensive growth renders the corresponding annealing protocol inaccessible to efficient Clifford-based simulation via the Gottesman–Knill theorem [103]. The improvement in the spectral gap delivered by the non-stoquastic protocol is thus correlated with an increased classical hardness of the instantaneous ground state, in agreement with the complexity-theoretic motivation for non-stoquasticity discussed in Sec. 5.3.

The p -spin model demonstrates the generation of substantial entanglement and SRE for non-stoquastic protocols in a permutation-symmetric setting where large system sizes are accessible to numerical simulation. We now ask whether the same observed correlations between gap structure and resource generation occur in a geometrically local Ising model without permutation symmetry.

6.3 Numerical results for the geometrically local Ising model

As a second benchmark problem, we study the geometrically local Ising model introduced by Albash in Ref. [294], for which a potential speedup with non-stoquastic annealing protocols has been reported. Unlike the p -spin model, this system lacks permutation symmetry, involves only two-body interactions, and features geometrically local qubit connectivity, all of which makes it more representative of near-term hardware implementations at the cost of restricting the analysis to smaller system sizes (see Sec. 3.2.2 for details).

After defining the non-stoquastic annealing protocol (Sec. 6.3.1), we use the global transverse magnetization to map out the phase structure induced by the catalyst (Sec. 6.3.2) and provide a microscopic interpretation of the resulting plateau structure (Sec. 6.3.3). We then analyze the energy gap, bipartite entanglement entropy, and second stabilizer Rényi entropy across the annealing parameter space (Sec. 6.3.4) and study their finite-size scaling at fixed catalyst strength (Sec. 6.3.5).

6.3.1 Non-stoquastic annealing protocol

We recall from Sec. 3.2.2 the structure of the model: N qubits arranged on two rings of $N/2$ sites each, connected by two inter-ring couplings linking diametrically opposite sites (see Fig. 3.2). The cost Hamiltonian

$$H_{\text{cost}} = \sum_{\langle i,j \rangle} J_{ij} \hat{\sigma}_i^z \hat{\sigma}_j^z \quad (6.12)$$

combines ferromagnetic and antiferromagnetic couplings of three different magnitudes along the edges of this graph, while the standard transverse-field driver reads $H_{\text{drive}} = -\sum_i \hat{\sigma}_i^x$.

Following Ref. [294], we introduce a non-stoquastic catalyst Hamiltonian sharing the same connectivity as H_{cost} ,

$$H_{\text{catalyst}} = \sum_{\langle i,j \rangle} \hat{\sigma}_i^x \hat{\sigma}_j^x, \quad (6.13)$$

and the total time-dependent annealing Hamiltonian:

$$\hat{H}(s, \lambda) = (1 - s) H_{\text{drive}} + s H_{\text{cost}} + \lambda s(1 - s) H_{\text{catalyst}}. \quad (6.14)$$

The schedule of Eq. (6.14) differs slightly from the protocol used for the p -spin model in Eq. (6.8): here λ is a fixed real parameter that controls the catalyst strength, while the envelope $s(1 - s)$ ensures that the catalyst is active only at intermediate times and vanishes at the start- and endpoints of the sweep. The Hamiltonian therefore reduces to H_{drive} at $s = 0$ and to H_{cost} at $s = 1$ for any value of λ , so that the boundary conditions of the protocol are independent of the catalyst strength.

For $\lambda > 0$, the catalyst Hamiltonian introduces positive off-diagonal matrix elements in the computational basis, rendering the total Hamiltonian non-stoquastic. The stoquastic regime $\lambda \leq 0$ has been studied in Ref. [294], where the minimum energy gap was shown to close exponentially with system size due to a perturbative level crossing near $s = 1$. For a suitable positive λ , the catalyst splits this crossing into multiple avoided crossings, leading to a significantly improved, potentially polynomial scaling of the gap.

The Hamiltonian (6.14) preserves the $\mathbb{Z}_2 \times \mathbb{Z}_2$ symmetry group generated by the global spin flip $\hat{\Pi} = \prod_{i=1}^N \hat{\sigma}_i^x$ and the ring-swap operation $\hat{P}_{\text{swap}}$ at every (s, λ) , so the dynamics initialized in $|+\rangle^{\otimes N}$ remains confined to the joint $+1$ eigenspace of both operators throughout the protocol. The construction of an orthonormal basis for this symmetric sector and the resulting constant-factor reduction of the Hilbert-space dimension to $d \sim 2^N/4$ were discussed in Sec. 3.2.2. The bipartite entanglement entropy is evaluated from the ground-state wave function by an exact Schmidt decomposition across a balanced bipartition. For the stabilizer Rényi entropy, whose exact evaluation requires summing over 4^N Pauli strings, we use Monte Carlo sampling over Pauli strings following Ref. [386]. This gives access to system sizes up to $N = 20$.

6.3.2 Transitions in the magnetization landscape

To characterize the behavior of the system along the annealing path, we study order parameters as functions of s and λ . Throughout this section, $|\psi_0(s, \lambda)\rangle$ denotes the instantaneous ground state of the Hamiltonian (6.14).

A natural diagnostic is the global transverse magnetization

$$\langle m_x \rangle = \frac{1}{N} \sum_{i=1}^N \langle \psi_0 | \hat{\sigma}_i^x | \psi_0 \rangle, \quad (6.15)$$

which equals unity at $s = 0$, where the ground state is $|+\rangle^{\otimes N}$, and vanishes at $s = 1$, where the ground state is an eigenstate of H_{cost} and therefore has no transverse coherence. The behavior of $\langle m_x \rangle$ along the sweep reflects how the ground state reorganizes along the annealing path: sharp variations in the thermodynamic limit signal a first-order transition, while smooth variations are characteristic of a second-order transition. We use this observable as our primary diagnostic to identify the different regimes induced by the catalyst.

Figure 6.3 shows the magnetization landscape for system sizes $N = 8, 10, 12, 14, 16$. The first column [panels (a1)–(e1)] reports heatmaps of $\langle m_x \rangle$ as a function of s and λ . For $\lambda < 0$, the magnetization varies smoothly across the annealing path, with no sharp features except near the final stage where the system approaches the ground state of the cost Hamiltonian. For $\lambda > 0$, a richer structure develops and becomes increasingly regular at large λ : regions of rapid variation emerge in the (s, λ) plane and sharpen with increasing system size, signaling

the onset of a phase transition in the thermodynamic limit. At large λ , the magnetization develops a sequence of plateaus at discrete values, whose number grows with N .

For finite N , the magnetization profile remains continuous: the apparent plateaus are connected by narrow but smooth rearrangements of the ground state that sharpen into steps only in the thermodynamic limit. These rearrangements originate from a sequence of avoided level crossings between competing low-energy configurations, as we make explicit in Sec. 6.3.3 below. The third column of Fig. 6.3 [panels (a3)–(e3)] shows $\langle m_x \rangle$ along a vertical sweep at fixed $\lambda = 4$, where the plateau structure and the sharp transitions between plateaus are directly visible.

The spatial structure of the plateau states is resolved by the local transverse magnetization,

$$\langle m_x^{(i)} \rangle = \langle \psi_0 | \hat{\sigma}_i^x | \psi_0 \rangle, \quad (6.16)$$

shown for $\lambda = 4$ in the second column of Fig. 6.3 [panels (a2)–(e2)]. Because the model lacks permutation symmetry, $\langle m_x^{(i)} \rangle$ varies across lattice sites, and each plateau is characterized by a distinct pattern in which specific sites retain transverse polarization while others get depolarized.

6.3.3 Microscopic description of the plateau structure

The plateau structure observed in the previous subsection admits a quantitative microscopic interpretation in the regime of large catalyst strength and early annealing time, where the cost Hamiltonian contributes only perturbatively. Within this effective description, each plateau corresponds to a classical spin configuration of an effective XX -type model, and the boundaries between plateaus follow from explicit level crossings between these configurations.

Effective Hamiltonian. We start from the full annealing Hamiltonian, Eq. (6.14), in the regime $\lambda s(1-s) \gg s$. Keeping the product $\tilde{\lambda} \equiv \lambda s$ fixed and expanding in s , the envelope $s(1-s)$ reduces to s at leading order and the effective Hamiltonian becomes

$$\hat{H}_{\text{eff}} = - \sum_i \hat{\sigma}_i^x + \tilde{\lambda} \sum_{\langle i,j \rangle} \hat{\sigma}_i^x \hat{\sigma}_j^x + \mathcal{O}(s). \quad (6.17)$$

At leading order the cost Hamiltonian drops out, and the dynamics is governed by a transverse-field model with competing on-site and nearest-neighbor XX couplings. All terms in Eq. (6.17) are diagonal in the $\hat{\sigma}^x$ eigenbasis and commute pairwise, so the ground state is the minimum-energy eigenstate over classical spin configurations $\{s_i\}$ with $s_i = \pm 1$, whose energy reads

$$E(\{s_i\}) = - \sum_i s_i + \tilde{\lambda} \sum_{\langle i,j \rangle} s_i s_j. \quad (6.18)$$

The competition between the two terms is controlled entirely by $\tilde{\lambda}$. For $\tilde{\lambda} \ll 1$ the on-site term dominates and the ground state is fully polarized along $+x$. For $\tilde{\lambda} \gg 1$ the XX term dominates and favors configurations with $s_i s_j = -1$ on as many bonds as possible, namely antiferromagnetic arrangements in the $\hat{\sigma}^x$ basis with vanishing or minimal transverse magnetization. At intermediate $\tilde{\lambda}$, neither term wins globally, and a discrete sequence of configurations becomes successively favored as $\tilde{\lambda}$ grows.

For fixed large catalyst strength λ , the parameter $\tilde{\lambda} = \lambda s$ should therefore be interpreted as a rescaled early-time coordinate. The regime $\tilde{\lambda} \ll 1$ corresponds to the very beginning




Figure 6.3: Magnetization across the phase diagrams and spin-resolved dynamics of the local Ising model [Eq. (6.14)] for system sizes $N = 8$ – 16 (rows). **Left column:** heatmaps of the global transverse magnetization $\langle m_x \rangle$ as a function of the annealing parameter s and the catalyst strength λ . **Middle column:** spin-resolved local transverse magnetization $\langle m_x^{(i)} \rangle$ versus s at fixed $\lambda = 4$. **Right column:** vertical annealing sweeps at fixed $\lambda = 4$, highlighting the step-like plateaus in $\langle m_x \rangle$ and the abrupt crossovers between magnetization sectors as s varies.

of the sweep, $s \ll 1/\lambda$, where the transverse field dominates. The plateau transitions occur at $s = O(1/\lambda)$, where $\tilde{\lambda} = O(1)$. Finally, $\tilde{\lambda} \gg 1$ corresponds to times later within the small- s window, $1/\lambda \ll s \ll 1$, but not to the end of the annealing sweep. In this regime the XX catalyst term dominates the effective Hamiltonian, while the cost Hamiltonian remains perturbative as long as $s = \tilde{\lambda}/\lambda \ll 1$.

Classical plateau configurations. Let $\mathcal{G}$ denote the two-ring graph defined by H_{cost} , with $|E(\mathcal{G})| = N + 2$ edges. The graph contains four *edge spins* (the two sites on each ring that participate in the inter-ring couplings) of coordination number three, and $N - 4$ *bulk spins* of coordination number two. Starting from the fully polarized configuration with all $s_i = +1$, a discrete family of configurations becomes energetically favored as $\tilde{\lambda}$ grows, each obtained by flipping a specific subset S of spins. These configurations correspond to the plateaus observed in Fig. 6.3 and are illustrated in Figs. 6.4 and 6.5 for $N = 10$ and $N = 12$, respectively.

For a configuration obtained from the fully polarized state by flipping a subset S of spins, the field contribution is changed only by the number of flipped spins. Since each flipped spin changes $s_i = +1$ to $s_i = -1$, the first term in Eq. (6.18) changes from $-N$ to $-N + 2|S|$. The interaction term is modified only on bonds connecting a flipped spin to an unflipped spin: on such a bond, $s_i s_j$ changes from $+1$ to -1 , lowering the XX -bond contribution by $2\tilde{\lambda}$. Bonds whose endpoints are both flipped, or both unflipped, remain unchanged. If $B(S)$ denotes the number of bonds with one endpoint in S and the other outside S , the energy then reads

$$E(S) = -[N - 2|S|] + \tilde{\lambda}[(N + 2) - 2B(S)]. \quad (6.19)$$

Thus, energetically favorable flipped-spin configurations are those that gain many antiferromagnetically aligned XX bonds while paying the energy cost associated with the flipped spins.

Two limiting cases are worth highlighting. The fully polarized initial state ($S = \emptyset$) has all bonds contributing $+\tilde{\lambda}$, yielding

$$E_{\text{pol}} = -N + \tilde{\lambda}(N + 2). \quad (6.20)$$

At the opposite end, for bipartite rings ($N/2$ even) all bonds can be simultaneously antiferromagnetic and the total magnetization vanishes, giving

$$E_{\text{AFM}} = -\tilde{\lambda}(N + 2). \quad (6.21)$$

For odd-length rings ($N/2$ odd, e.g., $N = 10$) full antiferromagnetic alignment is obstructed by frustration, so the actual final-state energy is slightly higher than Eq. (6.21).

The first intermediate plateau is obtained by flipping a pair of edge spins. Edge spins have higher coordination than bulk spins, so each contributes more strongly to the XX term when flipped, and is therefore favored to flip first as $\tilde{\lambda}$ increases. Among all pairs, the one minimizing the energy is the configuration with two edge spins that are not directly connected by an inter-ring bond [see Figs. 6.4(b) and 6.5(b)]. For this choice $|S| = 2$ and $B(S) = 6$, namely the 6 bonds incident to the two flipped edge spins, none of which lies between them, yielding

$$E_{\text{plateau},1} = -N + 4 + \tilde{\lambda}(N - 10). \quad (6.22)$$

Equation (6.22) is independent of which non-adjacent pair of edge spins is chosen, since the graph (and the effective Hamiltonian, Eq. (6.17)) is symmetric with respect to the four edge spins. The actual plateau state is a symmetric superposition over all representatives




Figure 6.4: Representative spin configurations of the plateau states of the local Ising model for $N = 10$, shown in the $\hat{\sigma}^x$ basis. (a) Fully polarized initial state, $s_i = +1$ for all i . (b) First intermediate plateau, with two non-adjacent edge spins flipped to $s_i = -1$. (c) Final plateau, characterized by strongly reduced transverse magnetization (full antiferromagnetic alignment is obstructed by frustration on the odd-length rings). Arrows indicate the classical value $s_i = \pm 1$ at each site. Only one representative configuration is shown per plateau. The true plateau state is a symmetric superposition of all configurations (sharing the same energy) related by the spatial symmetries of the effective Hamiltonian (Eq. (6.17)).

related by this spatial symmetry, each sharing the same classical energy. This superposition is responsible for the depolarization of the four edge spins observed in the spin-resolved magnetization profiles of Fig. 6.3: while each individual representative has fully polarized edge spins with $s_i = \pm 1$, their symmetric superposition has $\langle \hat{\sigma}_i^x \rangle = 0$ at every edge site, while the bulk spins, which are equally polarized in every representative configuration, also remain polarized in the superposition.

We note that, since the protocol is initialized in the $+1$ eigensector of $\hat{\Pi}$ and parity is conserved by the dynamics, only configurations with an even number of flipped spins are accessible, so the plateau sequence proceeds in steps of two flipped spins. For $N = 8$ and $N = 10$, Eq. (6.22) describes the unique intermediate configuration between the polarized and the final states. For $N > 10$, additional plateaus appear, obtained by flipping further pairs of spins while keeping $|\partial S|$ as large as possible. As a representative example, for $N = 12$ the second intermediate plateau [Fig. 6.5(c)] has $|S| = 4$ and $B(S) = 10$, giving

$$E_{\text{plateau},2}^{(N=12)} = -4 - 6\tilde{\lambda}. \quad (6.23)$$

Plateau boundaries. The boundaries between plateaus in the (λ, s) plane follow from the level crossings of the classical energies above. Equating E_{pol} with $E_{\text{plateau},1}$ yields, independently of N ,

$$\tilde{\lambda}_1 = \frac{1}{3} \iff s_1 = \frac{1}{3\lambda}, \quad (6.24)$$

which marks the transition out of the fully polarized configuration. Equating $E_{\text{plateau},1}$ with E_{AFM} further yields

$$\tilde{\lambda}_2 = \frac{1}{2} \iff s_2 = \frac{1}{2\lambda}, \quad (6.25)$$




Figure 6.5: Same as Fig. 6.4, but for $N = 12$. The longer bipartite rings support an additional intermediate configuration between the first plateau and the final state.




Figure 6.6: Low-lying energy spectrum of the local Ising model with $N = 12$ and $\lambda = 4$ as a function of the normalized annealing time $s \in [0.05, 0.2]$. Grey curves show the numerically exact spectrum, with the ground state highlighted in dark grey. Colored lines show the analytical plateau energies of Eqs. (6.20), (6.22), (6.23), and (6.21). Red circles mark the three avoided crossings at $s \approx 0.083, 0.116, 0.136$.

which for $N \in \{8, 10\}$ identifies the transition time to the final configuration directly. For $N = 12$, inserting Eq. (6.23) one finds that both $E_{\text{plateau},1} = E_{\text{plateau},2}$ and $E_{\text{plateau},2} = E_{\text{AFM}}$ are satisfied at the same value $\tilde{\lambda} = 1/2$. Within the strictly classical analysis, the second plateau is therefore degenerate with the first plateau and the final state at a single point and is never the unique minimum.

Comparison with the exact spectrum. Figure 6.6 compares the analytical predictions with the numerically exact low-lying spectrum for $N = 12$ and $\lambda = 4$, in the window $s \in [0.05, 0.2]$ where all relevant rearrangements take place. Three avoided crossings, marked by red circles at $s \approx 0.083, 0.116, 0.136$, separate four windows of the spectrum whose ground-state slopes agree with those of the analytical predictions in the order prescribed

by the analysis above. The first avoided crossing is centered close to the analytical value $s_1 = 1/(3\lambda) \approx 0.083$, while the remaining two lie symmetrically around the degenerate classical value $s_2 = 1/(2\lambda) = 0.125$.

Two systematic deviations from the strictly classical analysis are visible. First, the exact ground-state energy lies above the analytical predictions by an offset that depends weakly on s . This offset reflects the $\mathcal{O}(s)$ contributions neglected in Eq. (6.17): at $\lambda = 4$ and $s \sim 0.1$, the prefactor in front of the driving Hamiltonian is $(1 - s) \neq 1$ and also the cost Hamiltonian is not entirely negligible; the resulting effect is that the energies are shifted upwards without strongly altering their slopes within the observed window. Second, the second-plateau configuration emerges as the exact ground state over a finite window between the second and third avoided crossings, despite being merely degenerate at the level of the classical analysis. The lifting of this degeneracy is again a consequence of the sub-leading cost Hamiltonian contributions, which select the second-plateau configuration in this window.

Apart from these two effects, the analytical plateau structure accurately captures the location and ordering of the transitions, with the classical prediction becoming asymptotically exact in the large- λ limit. The phase diagram observed in Fig. 6.3 therefore admits, in the large- λ , small- s regime, a quantitative microscopic interpretation in terms of competing classical configurations of the effective Hamiltonian Eq. (6.17), with the full Hamiltonian providing systematic corrections that resolve the outcome at the classical degeneracy points.

This picture clarifies the role of the non-stoquastic catalyst. Already at small s , the XX term drives the ground state through a sequence of rearrangements with progressively smaller global magnetization, rather than through a single transition. These intermediate rearrangements do not by themselves guarantee improved annealing performance, which arises instead from the interplay between cost and catalyst Hamiltonian and is reported in Ref. [294] to be optimal around $\lambda \sim 1$ (with a strong dependence on system size).

6.3.4 Quantum resources across the annealing parameter space

As in the case of the p -spin model (Sec. 6.2.2), we now characterize the system by three observables evaluated on the instantaneous ground state: the energy gap, the bipartite entanglement entropy, and the second stabilizer Rényi entropy. We consider these observables as a function of the annealing parameters (λ, s) on a 101×401 grid covering $\lambda \in [-2, 2]$ and $s \in [0, 1]$. Figure 6.7 reports the resulting heatmaps for system sizes $N = 8, 10, 12, 14, 16$, with the energy gap, the entanglement entropy, and the SRE density shown in the first, second, and third columns, respectively.

For $\lambda \leq 0$, the three observables exhibit local extrema around the same single dominant line in the (λ, s) plane: along this line, the energy gap and the entanglement entropy develop a pronounced minimum and maximum, respectively, while the SRE exhibits a corresponding local maximum. The features become increasingly narrow as the system size grows, indicating that in the stoquastic regime the annealing dynamics is dominated by a single rearrangement of the ground state.

For $\lambda > 0$, the structure of the annealing landscape changes qualitatively, in agreement with the magnetization data of Fig. 6.3. In this regime, the three observables instead display multiple, broader local extrema. In particular, the energy gap develops several local minima, each accompanied by corresponding local maxima of the entanglement entropy and the SRE. This pattern reflects the multiple competing rearrangements of the ground state induced by the non-stoquastic catalyst, as analyzed microscopically in Sec. 6.3.3: each plateau boundary appears as a local minimum of the gap, accompanied by a local resource peak.




Figure 6.7: **First column:** energy gap as a function of the annealing parameters s and λ , for system sizes $N = 8, 10, 12, 14, 16$ (from top to bottom). **Middle column:** bipartite entanglement entropy evaluated on the instantaneous ground state of the Hamiltonian in Eq. (6.14) for the same parameter range and system sizes. **Third column:** stabilizer Rényi entropy density, computed under the same conditions.

The same qualitative correlation between spectral bottleneck and quantum-resource generation observed for the p -spin model therefore survives the loss of permutation symmetry and the presence of a sequence of spin rearrangements. The resource peaks, however, are now distributed over multiple values of s rather than concentrated at a single transition, mirroring the structure of the underlying magnetization plateaus.

6.3.5 Finite-size scaling along annealing protocols

To quantify the scaling of the quantum resources with system size, we now consider annealing protocols at fixed values of λ and extract, along each sweep, the minimum energy gap and

the maximum values of the entanglement entropy and the stabilizer Rényi entropy. The resulting data are shown in Fig. 6.8 as a function of N .

The annealing parameter s is discretized non-uniformly, with observable- and system-size-dependent resolutions chosen to accurately resolve the relevant extrema. For the stoquastic case ($\lambda = 0$), all observables are evaluated on a uniform grid with $\Delta s = 10^{-4}$. For the non-stoquastic case ($\lambda = 4$), finer resolutions are used in restricted intervals around the extrema: $\Delta s = 5 \times 10^{-5}$ for the energy gap at $N = 8, 10$ and $\Delta s = 10^{-5}$ for larger sizes; analogous resolutions are adopted for the SRE, except for $N = 18, 20$ where $\Delta s = 5 \times 10^{-5}$ is used; and $\Delta s = 10^{-5}$ for the entanglement entropy at all available system sizes.

We focus on two representative cases: the stoquastic protocol with $\lambda = 0$ and a non-stoquastic protocol with $\lambda = 4$. Smaller positive values of λ feature a stronger competition between the cost and catalyst Hamiltonians, leading to a more intricate energy landscape and obscuring the scaling behavior, while $\lambda = 4$ yields more regular structures that allow for a more systematic analysis. We emphasize, however, that no genuine non-stoquastic advantage is expected at $\lambda = 4$, which lies far from the conjectured optimum around $\lambda \sim 1$ [294]. To account for the multiple local extrema present at $\lambda = 4$, we report both the global minimum and selected local minima for the energy gap, namely the first and the last minimum (Fig. 6.8 (a)). For the entanglement entropy and the SRE, the structure during the annealing protocol is more complicated, with both local extrema and plateau-like regions that are not always clearly separable from each other. We thus only report the maximum value attained during the sweep (Fig. 6.8 (b) and (c)).

For the stoquastic protocol ($\lambda = 0$), the minimum gap decreases exponentially with system size, in agreement with Ref. [294]. The entanglement entropy remains approximately constant in this regime, while the stabilizer Rényi entropy grows roughly linearly with N , indicating an extensive generation of non-stabilizerness already in the absence of the catalyst.

For the non-stoquastic protocol, the behavior is more complex. Both the first (corresponding to the global) and the last energy gap minima show a slower decrease compared to the stoquastic case. Although the precise functional form cannot be unambiguously determined from the accessible system sizes, the available data suggest an improvement in the scaling and thus in the possible annealing speed. The entanglement entropy exhibits larger values and a more pronounced growth with system size compared to the stoquastic case, reflecting enhanced quantum correlations induced by the catalyst.

The behavior of the stabilizer Rényi entropy is similar to the stoquastic case: when considering the global maximum along the annealing path, the SRE grows approximately linearly with N , indicating an overall extensive scaling also in the non-stoquastic case.

Overall, these finite-size results indicate that the non-stoquastic catalyst does more than modify the minimum spectral gap: it reorganizes the structure of the instantaneous ground state and redistributes and induces the generation of quantum resources along the annealing path. In the p-spin model, the regimes in which first-order transitions are avoided coincide with enhanced entanglement and substantially larger extensive stabilizer Rényi entropy, suggesting that the improved gap scaling is accompanied by states that are harder to describe by both tensor-network and stabilizer-based methods. In the geometrically local Ising model, the catalyst splits the dominant rearrangement of the stoquastic protocol into a sequence of plateau transitions, producing multiple spectral bottlenecks and a correspondingly richer resource landscape. Although the accessible system sizes do not allow a definitive asymptotic scaling analysis, the results support the view that non-stoquastic catalysts can improve annealing performance only by generating nontrivial many-body structure along the path (at least for the in this chapter considered models). Extending this analysis to different systems, optimized catalyst strengths, and more experimentally constrained protocols will




Figure 6.8: Finite-size scaling of (a) the minimum energy gap, (b) the maximum bipartite entanglement entropy, and (c) the maximum stabilizer Rényi entropy along annealing protocols with fixed values of λ . Results are shown for the stoquastic case $\lambda = 0$ (blue squares with solid lines) and the non-stoquastic case $\lambda = 4$. In the non-stoquastic case, several local minima of the energy gap occur along the annealing path. We report the absolute minimum (red dots with solid lines), as well as the first (yellow, dashed lines) and last (green, dashed lines) local minima. For the entanglement entropy and SRE, whose profiles exhibit a more complicated structure without clearly separable peaks and plateaus, we report only the maximum value attained along the path. Lines are intended as guides to the eye.

be essential for determining whether these correlated enhancements of gap, entanglement, and non-stabilizerness can translate into a practical quantum advantage.

Conclusions and Outlook

In this thesis, we have addressed the question of how quantum annealing can be brought closer to practical relevance on near-term trapped-ion hardware, approaching it from two complementary angles. On the hardware side, we have analyzed the structure of the magnetic-gradient-induced coupling (MAGIC) scheme for trapped-ion quantum computing beyond its standard leading-order description, and we have designed a noise mitigation protocol tailored to annealing algorithms on this platform. On the algorithmic side, we have investigated how the structure of the annealing Hamiltonian itself can be modified to enlarge the minimum spectral gap that controls the runtime of annealing protocols. Furthermore, we studied how quantum resources behave for stoquastic and non-stoquastic annealing schedules and how they relate to the difficulty of classical simulation. In this chapter, we summarize the main results along these two lines and outline directions for future research.

Summary of main results

Higher-order corrections to the MAGIC scheme. In Chapter 2, based on Ref. [205], we have systematically analyzed the structure and effects of higher-order terms in the MAGIC setup arising from anharmonicities of the external potential, including Coulomb repulsion between ions and curvature of the applied external magnetic field gradient. After following the standard derivation [161], decoupling spin and phonon degrees of freedom at leading order, the higher-order terms generate three-body spin couplings, additional two-body and local-field contributions, and various phonon-phonon conversion mechanisms. By computing their natural strength using experimentally realistic parameters [163], we have shown that most of these contributions are negligible in realistic situations. Two terms, however, can have significant impact and need to be considered. The first is a phonon-occupation-dependent longitudinal field whose strength increases with chain length and phonon excitation numbers, and which can become comparable to the native two-body coupling strength when scaling to larger chains. While its mean effect can be compensated through recalibration, virtual Z rotations or microwave drive detuning, phonon number fluctuations remain problematic and underline the importance of efficient cooling techniques such as sympathetic cooling. The second is a two-to-one phonon conversion process that is suppressed under ground-state cooling and can additionally be avoided through suitable trap design that detunes phonon frequencies away from resonance. Importantly, our analysis also shows that the small native three-body interaction strengths produced by these higher-order corrections, which are orders of magnitude smaller than the native spin-spin couplings for typical parameters, are not by themselves sufficient to be exploited as a resource for quantum information processing. We have placed this finding in context by comparing it to alternative schemes for engineering three-body interactions in trapped ions.

Dynamical decoupling for analog quantum annealing. In Chapter 4, based on Ref. [144], we have addressed the most prominent error channel for magnetically sensitive trapped-ion qubits, namely fluctuations of external magnetic fields, which translate into fluctuating longitudinal fields on the qubits. By analyzing how local-field disorder of realistic amplitude affects the ground state of the cost Hamiltonian in quantum annealing, we have shown that moderate disorder strengths are sufficient to alter the encoded solution and thus prevent successful annealing without active noise suppression. To mitigate this effect, we have proposed an annealing protocol in which the optimization problem is encoded exclusively in two-body interactions through a single ancilla qubit, supplemented by periodic global spin-flip pulses interleaved into the annealing schedule. The ancilla encoding ensures that the dynamical decoupling pulses commute with the cost Hamiltonian, so that the encoded ground state is preserved while the longitudinal noise is averaged out. Through a Magnus expansion of the time-evolution operator, we have shown analytically that the leading-order corrections to the ideal annealing dynamics are controlled by the product of noise amplitude and the time interval between consecutive pulses. Numerical simulations on industrially relevant benchmarks, namely minimal Multiple Object Tracking instances of five and nine qubits, and cutting-stock instances of five and six qubits, as well as Sherrington-Kirkpatrick instances up to twelve qubits have confirmed that pulse rates of approximately 2.5 ms^{-1} are sufficient to restore nearly noise-free fidelity across realistic noise spectra and amplitudes up to a worst-case 1 kHz, well within the capabilities of current trapped-ion hardware. We have further identified a universal scaling behavior of the final fidelity in terms of a generalized parameter that combines noise amplitude and pulse interval, providing a compact figure of merit for designing decoupling sequences across different noise environments. Although our analysis was grounded in the MAGIC platform, the protocol and its scaling behavior apply broadly to other quantum computing architectures that suffer from longitudinal-field noise.

Direct PUBO encoding of optimization problems. In Chapter 5, based on Ref. [203], we have studied the design of the cost Hamiltonian itself, asking whether the standard practice of reducing higher-order optimization problems to quadratic unconstrained binary optimization (QUBO) form is justified, or whether retaining the native polynomial unconstrained binary optimization (PUBO) structure can be advantageous for some problems. We have shown that direct PUBO formulations can lead to substantial savings in the number of qubits, since the QUBO reduction typically requires ancillary slack variables and energy penalties to enforce constraints. More importantly, by exact diagonalization on small-scale hard-to-solve 3-SAT instances generated by different benchmark generators, we have shown that the minimum energy gap during the annealing sweep can be significantly larger in the PUBO formulation than in the equivalent QUBO reduction. For 3-SAT instances generated by `toughSAT` (random instances postselected for unique solutions), we have observed a significant improvement in the gap-scaling exponent, which translates into an exponentially faster adiabatic sweep, compared to the equivalent QUBO reduction, with our estimates suggesting a speedup of about two orders of magnitude already for problems with only eleven classical variables. For instances generated by `uniquePT1` (structured hard-to-solve benchmark problem instances), the scaling exponent is comparable in PUBO and QUBO, but the larger absolute gap still yields a constant-factor speedup. We have shown that this advantage persists even when accounting for the overhead of synthesizing three-qubit interactions from universal sets of single- and two-qubit gates on digital architectures, where four CNOT gates suffice to implement the required $\sigma^z \sigma^z \sigma^z$ rotation. As a side result, our analysis of different 3-SAT generators has revealed that a third generator, `uniquePT4`, produces only

trivially solvable instances whose PUBO cost Hamiltonians reduce to sums of local fields, a finding relevant for benchmarking both quantum and classical solvers.

Non-stoquastic annealing and quantum resources. In Chapter 6, based on Ref. [204], we have addressed the question of whether the spectral improvements obtained through non-stoquastic catalysts are accompanied by the generation of quantum resources that make the resulting states genuinely hard to simulate classically. To this end, we have monitored two complementary resources of the instantaneous ground state along the annealing path, the bipartite entanglement entropy and the second stabilizer Rényi entropy (as a measure for non-stabilizerness), which respectively quantify the hardness of classical simulation by tensor-network methods and by stabilizer-tableau methods. We have analyzed two benchmark models, the fully connected ferromagnetic p -spin model and a geometrically local Ising model. In the p -spin model, the entanglement entropy of the instantaneous ground state saturates with system size while the stabilizer Rényi entropy grows linearly, with a slope that increases sharply when the annealing path crosses from a first-order to a second-order phase transition under the action of the catalyst. In the local Ising model, the catalyst induces a sequence of avoided crossings, each accompanied by local maxima or plateaus of both quantum resources, and we have provided a microscopic interpretation of these crossings in terms of competing classical configurations of an effective Hamiltonian valid in the regime of large catalyst strength and early annealing time. In both models, the regimes in which non-stoquasticity improves the minimum spectral gap coincide with those in which the ground state develops larger values of both quantum resources. This concurrent behavior identifies these non-stoquastic annealing protocols as candidates for regimes in which a potential efficient quantum algorithm coincides with classical hardness.

Outlook and future directions

The work presented in this thesis suggests several natural directions for future research.

The analysis of perturbative corrections in the MAGIC scheme can be extended to higher-order terms in the Coulomb potential and to a more detailed treatment of the radial phonon modes, which become relevant when scaling to larger ion chains. Similar analyses may also prove valuable for related platforms based on Mølmer-Sørensen-type gates or oscillating magnetic fields, as well as for superconducting qubit architectures governed by similar microscopic Hamiltonians. The dynamical decoupling protocol developed here can in turn be extended to include more sophisticated decoupling sequences, such as universally robust pulse sequences that self-correct against control errors, and a broader spectrum of noise sources. A particularly intriguing theoretical question is whether the universal scaling behavior we have identified, in which the final fidelity depends on a simple combination of noise amplitude and pulse interval, can be derived analytically from the noise spectral density, which would provide a guiding principle for designing optimal decoupling sequences in arbitrary noise environments. On the experimental side, the protocol can also serve as a controllable knob to tune the impact of decoherence during annealing, providing a means to systematically study the role of quantum effects in optimization performance.

Our finding that direct PUBO formulations can exponentially improve the gap-scaling exponent for specific classes of 3-SAT instances raises the question of how broadly this advantage extends. Systematic studies on other natively higher-order problems, together with careful comparisons against state-of-the-art classical solvers, would help identify regimes in which quantum annealing with PUBO encodings is competitive with or surpasses classi-

cal performance. The connection between non-stoquastic catalysts and quantum resources identified here similarly invites investigation into larger problem classes without exploitable symmetries, including random spin glasses and industrially motivated optimization problems. The microscopic interpretation we have provided for the geometrically local Ising model, in terms of competing classical configurations of an effective Hamiltonian, may also serve as a starting point for the systematic design of catalyst Hamiltonians tailored to specific problem structures.

A natural next step beyond the individual directions outlined above is their integration. A trapped-ion experiment combining a well-characterized hardware platform, active noise mitigation through dynamical decoupling, and algorithmic ingredients such as native PUBO encodings or non-stoquastic catalysts would unite the threads developed in this thesis into a single setting, and would constitute a meaningful step from current proof-of-principle demonstrations toward devices capable of tackling optimization problems of practical interest.

Bibliography

- [1] J. P. Dowling and G. J. Milburn, “Quantum technology: the second quantum revolution”, *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* **361**, edited by A. G. J. MacFarlane, 1655–1674 (2003).
- [2] A. Acín, I. Bloch, H. Buhrman, T. Calarco, C. Eichler, J. Eisert, D. Esteve, N. Gisin, S. J. Glaser, F. Jelezko, S. Kuhr, M. Lewenstein, M. F. Riedel, P. O. Schmidt, R. Thew, A. Wallraff, I. Walmsley, and F. K. Wilhelm, “The quantum technologies roadmap: a european community view”, *New Journal of Physics* **20**, 080201 (2018).
- [3] C. Couteau and S. Lazić, “How to define quantum technology”, *Nature Reviews Physics* **7**, 234–235 (2025).
- [4] C. L. Degen, F. Reinhard, and P. Cappellaro, “Quantum sensing”, *Rev. Mod. Phys.* **89**, 035002 (2017).
- [5] N. Gisin and R. Thew, “Quantum communication”, *Nature Photonics* **1**, 165–171 (2007).
- [6] K. Azuma, S. E. Economou, D. Elkouss, P. Hilaire, L. Jiang, H.-K. Lo, and I. Tzitrin, “Quantum repeaters: from quantum networks to the quantum internet”, *Rev. Mod. Phys.* **95**, 045006 (2023).
- [7] J. Preskill, *Quantum computing 40 years later*, 2021.
- [8] T. Hoeffler, T. Haener, and M. Troyer, “Disentangling hype from practicality: on realistically achieving quantum advantage”, 10.48550/ARXIV.2307.00523 (2023).
- [9] J. Eisert and J. Preskill, *Mind the gaps: the fraught road to quantum advantage*, 2025.
- [10] R. Babbush, R. King, S. Boixo, W. Huggins, T. Khattar, G. H. Low, J. R. McClean, T. O’Brien, and N. C. Rubin, *The grand challenge of quantum applications*, 2025.
- [11] J. Preskill, “Quantum computing in the NISQ era and beyond”, *Quantum* **2**, 79 (2018).
- [12] K. Bharti, A. Cervera-Lierta, T. H. Kyaw, T. Haug, S. Alperin-Lea, A. Anand, M. Degroote, H. Heimonen, J. S. Kottmann, T. Menke, W.-K. Mok, S. Sim, L.-C. Kwek, and A. Aspuru-Guzik, “Noisy intermediate-scale quantum algorithms”, *Reviews of Modern Physics* **94**, 015004 (2022).
- [13] P. Benioff, “The computer as a physical system: a microscopic quantum mechanical hamiltonian model of computers as represented by turing machines”, *Journal of Statistical Physics* **22**, 563–591 (1980).
- [14] R. P. Feynman, “Simulating physics with computers”, *International Journal of Theoretical Physics* **21**, 467–488 (1982).

- [15] D. Deutsch, “Quantum theory, the church–turing principle and the universal quantum computer”, Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences **400**, 97–117 (1985).
- [16] D. E. Deutsch, “Quantum computational networks”, Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences **425**, 73–90 (1989).
- [17] A. Chi-Chih Yao, “Quantum circuit complexity”, in Proceedings of 1993 IEEE 34th annual foundations of computer science, SFCS-93 (), pp. 352–361.
- [18] D. Deutsch and R. Jozsa, “Rapid solution of problems by quantum computation”, Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences **439**, 553–558 (1992).
- [19] E. Bernstein and U. Vazirani, “Quantum complexity theory”, SIAM Journal on Computing **26**, 1411–1473 (1997).
- [20] D. R. Simon, “On the power of quantum computation”, SIAM Journal on Computing **26**, 1474–1483 (1997).
- [21] P. W. Shor, “Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer”, SIAM Journal on Computing **26**, 1484–1509 (1997).
- [22] L. K. Grover, “A fast quantum mechanical algorithm for database search”, in Proceedings of the twenty-eighth annual ACM symposium on theory of computing - STOC ’96, STOC ’96 (1996), pp. 212–219.
- [23] P. W. Shor, “Scheme for reducing decoherence in quantum computer memory”, Physical Review A **52**, R2493–R2496 (1995).
- [24] A. M. Steane, “Error correcting codes in quantum theory”, Physical Review Letters **77**, 793–797 (1996).
- [25] A. R. Calderbank and P. W. Shor, “Good quantum error-correcting codes exist”, Physical Review A **54**, 1098–1105 (1996).
- [26] P. Shor, “Fault-tolerant quantum computation”, in Proceedings of 37th conference on foundations of computer science, SFCS-96 (), pp. 56–65.
- [27] E. Knill, R. Laflamme, and W. H. Zurek, “Resilient quantum computation: error models and thresholds”, Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences **454**, 365–384 (1998).
- [28] D. Aharonov and M. Ben-Or, “Fault-tolerant quantum computation with constant error rate”, SIAM Journal on Computing **38**, 1207–1282 (2008).
- [29] A. Y. Kitaev, “Quantum computations: algorithms and error correction”, Russian Mathematical Surveys **52**, 1191–1249 (1997).
- [30] E. Farhi, J. Goldstone, S. Gutmann, and M. Sipser, *Quantum computation by adiabatic evolution*, 2000.
- [31] R. Raussendorf and H. J. Briegel, “A one-way quantum computer”, Physical Review Letters **86**, 5188–5191 (2001).
- [32] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev, “Adiabatic quantum computation is equivalent to standard quantum computation”, SIAM Review **50**, 755 (2008).
- [33] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information: 10th anniversary edition* (Cambridge University Press, 2010).

- [34] G. Brassard, P. Hoyer, M. Mosca, and A. Tapp, “Quantum amplitude amplification and estimation”, 10.48550/ARXIV.QUANT-PH/0005055 (2000).
- [35] C. H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W. K. Wootters, “Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels”, *Physical Review Letters* **70**, 1895–1899 (1993).
- [36] A. Y. Kitaev, A. H. Shen, and M. N. Vyalyi, *Classical and quantum computation*, Vol. 47, Graduate Studies in Mathematics (American Mathematical Society, Providence, RI, 2002).
- [37] J. Kempe, A. Kitaev, and O. Regev, “The complexity of the local hamiltonian problem”, 10.48550/ARXIV.QUANT-PH/0406180 (2004).
- [38] R. Jozsa, “Quantum factoring, discrete logarithms, and the hidden subgroup problem”, *Computing in Science & Engineering* **3**, 34–43 (2001).
- [39] C. H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani, “Strengths and weaknesses of quantum computing”, *SIAM Journal on Computing* **26**, 1510–1523 (1997).
- [40] T. Yamakawa and M. Zhandry, “Verifiable quantum advantage without structure”, in 2022 IEEE 63rd annual symposium on foundations of computer science (FOCS) (Oct. 2022), pp. 69–74.
- [41] S. P. Jordan, N. Shutty, M. Wootters, A. Zalcman, A. Schmidhuber, R. King, S. V. Isakov, T. Khattar, and R. Babbush, “Optimization by decoded quantum interferometry”, *Nature* **646**, 831–836 (2025).
- [42] M. Cain, Q. Xu, R. King, L. R. B. Picard, H. Levine, M. Endres, J. Preskill, H.-Y. Huang, and D. Bluvstein, *Shor’s algorithm is possible with as few as 10,000 reconfigurable atomic qubits*, 2026.
- [43] R. Babbush, A. Zalcman, C. Gidney, M. Broughton, T. Khattar, H. Neven, T. Bergamaschi, J. Drake, and D. Boneh, *Securing elliptic curve cryptocurrencies against quantum vulnerabilities: resource estimates and mitigations*, 2026.
- [44] S. Lloyd, “Universal quantum simulators”, *Science* **273**, 1073–1078 (1996).
- [45] D. W. Berry, A. M. Childs, R. Cleve, R. Kothari, and R. D. Somma, “Simulating Hamiltonian dynamics with a truncated Taylor series”, *Physical Review Letters* **114**, 090502 (2015).
- [46] G. H. Low and I. L. Chuang, “Hamiltonian simulation by qubitization”, *Quantum* **3**, 163 (2019).
- [47] I. M. Georgescu, S. Ashhab, and F. Nori, “Quantum simulation”, *Reviews of Modern Physics* **86**, 153–185 (2014).
- [48] S. McArdle, S. Endo, A. Aspuru-Guzik, S. C. Benjamin, and X. Yuan, “Quantum computational chemistry”, *Reviews of Modern Physics* **92**, 015003 (2020).
- [49] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, and A. Kandala, “Evidence for the utility of quantum computing before fault tolerance”, *Nature* **618**, 500–505 (2023).
- [50] J. Tindall, M. Fishman, E. M. Stoudenmire, and D. Sels, “Efficient tensor network simulation of IBM’s eagle kicked Ising experiment”, *PRX Quantum* **5**, 010308 (2024).
- [51] F. Alam et al., *Programmable digital quantum simulation of 2d Fermi-Hubbard dynamics using 72 superconducting qubits*, 2025.

- [52] F. Alam et al., *Fermionic dynamics on a trapped-ion quantum computer beyond exact classical simulation*, 2025.
- [53] L. Vilchez-Estevez, R. A. Santos, S. Y. Wang, and F. M. Gambetta, “Extracting the spin excitation spectrum of a fermionic system using a quantum processor”, *Physical Review B* **112**, 10.1103/ydfw-k83n (2025).
- [54] E. Granet et al., *Superconducting pairing correlations on a trapped-ion quantum computer*, 2025.
- [55] T. A. Chowdhury, V. Korepin, V. R. Pascuzzi, and K. Yu, “Quantum utility in simulating the real-time dynamics of the Fermi–Hubbard model using superconducting quantum computers”, *Applied Physics Reviews* **13**, 10.1063/5.0306069 (2026).
- [56] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, “Quantum machine learning”, *Nature* **549**, 195–202 (2017).
- [57] P. Wittek, *Quantum machine learning* (Elsevier, 2014).
- [58] S. Lloyd, M. Mohseni, and P. Rebentrost, *Quantum algorithms for supervised and unsupervised machine learning*, Nov. 2013.
- [59] M. Schuld, R. Sweke, and J. J. Meyer, “Effect of data encoding on the expressive power of variational quantum-machine-learning models”, *Physical Review A* **103**, 032430 (2021).
- [60] D. Pastorello, *Concise guide to quantum machine learning* (Springer Nature Singapore, 2023).
- [61] E. Farhi, J. Goldstone, S. Gutmann, J. Lapan, A. Lundgren, and D. Preda, “A quantum adiabatic evolution algorithm applied to random instances of an np-complete problem”, *Science* **292**, 472–475 (2001).
- [62] A. Das and B. K. Chakrabarti, “Colloquium: quantum annealing and analog quantum computation”, *Reviews of Modern Physics* **80**, 1061–1081 (2008).
- [63] E. Farhi, J. Goldstone, and S. Gutmann, *A quantum approximate optimization algorithm*, 2014.
- [64] P. Hauke, H. G. Katzgraber, W. Lechner, H. Nishimori, and W. D. Oliver, “Perspectives of quantum annealing: methods and implementations”, *Reports on Progress in Physics* **83**, 054401 (2020).
- [65] K. Blekos, D. Brand, A. Ceschini, C.-H. Chou, R.-H. Li, K. Pandya, and A. Summer, “A review on quantum approximate optimization algorithm and its variants”, *Physics Reports* **1068**, 1–66 (2024).
- [66] T. F. Rønnow, Z. Wang, J. Job, S. Boixo, S. V. Isakov, D. Wecker, J. M. Martinis, D. A. Lidar, and M. Troyer, “Defining and detecting quantum speedup”, *Science* **345**, 420–424 (2014).
- [67] B. Barak and K. Marwaha, “Classical algorithms and quantum limitations for maximum cut on high-girth graphs”, en, in , Vol. 215 (2022), 14:1–14:21.
- [68] A. W. Harrow, A. Hassidim, and S. Lloyd, “Quantum algorithm for linear systems of equations”, *Physical Review Letters* **103**, 150502 (2009).
- [69] P. Rebentrost, M. Mohseni, and S. Lloyd, “Quantum support vector machine for big data classification”, *Physical Review Letters* **113**, 130503 (2014).
- [70] S. Aaronson, “Read the fine print”, *Nature Physics* **11**, 291–293 (2015).

-
- [71] E. Tang, “A quantum-inspired classical algorithm for recommendation systems”, in Proceedings of the 51st annual acm sigact symposium on theory of computing, STOC 2019 (2019), pp. 217–228.
 - [72] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, “Barren plateaus in quantum neural network training landscapes”, Nature Communications **9**, 10.1038/s41467-018-07090-4 (2018).
 - [73] M. Cerezo, A. Arrasmith, R. Babbush, S. C. Benjamin, S. Endo, K. Fujii, J. R. McClean, K. Mitarai, X. Yuan, L. Cincio, and P. J. Coles, “Variational quantum algorithms”, Nature Reviews Physics **3**, 625–644 (2021).
 - [74] A. Montanaro, “Quantum speedup of monte carlo methods”, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences **471**, 20150301 (2015).
 - [75] D. Rattacaso, D. Jaschke, A. Trovato, I. Siloi, and S. Montangero, “Quantum algorithms for compact polymer thermodynamics”, 10.48550/ARXIV.2603.12334 (2026).
 - [76] C. Micheletti, P. Hauke, and P. Faccioli, “Polymer physics by quantum computing”, Physical Review Letters **127**, 080501 (2021).
 - [77] F. Slongo, P. Hauke, P. Faccioli, and C. Micheletti, “Quantum-inspired encoding enhances stochastic sampling of soft matter systems”, Science Advances **9**, 10.1126/sciadv.adi0204 (2023).
 - [78] V. Panizza, A. Roggero, P. Hauke, and P. Faccioli, “Statistical mechanics of heteropolymers from lattice gauge theory”, Physical Review Letters **134**, 158101 (2025).
 - [79] P. Hauke, G. Mattiotti, and P. Faccioli, “Dominant reaction pathways by quantum computing”, Physical Review Letters **126**, 028104 (2021).
 - [80] D. Ghamari, P. Hauke, R. Covino, and P. Faccioli, “Sampling rare conformational transitions with a quantum computer”, Scientific Reports **12**, 10.1038/s41598-022-20032-x (2022).
 - [81] S. Aaronson and S.-H. Hung, “Certified randomness from quantum supremacy”, in Proceedings of the 55th annual acm symposium on theory of computing, STOC ’23 (2023), pp. 933–944.
 - [82] M. Liu et al., “Certified randomness using a trapped-ion quantum processor”, Nature **640**, 343–348 (2025).
 - [83] S. Aaronson, *Summer 2022 quantum supremacy updates*, <https://scottaaronson.blog/?p=6645>, Shtetl-Optimized blog post, August 2022.
 - [84] S. Aaronson, “How much structure is needed for huge quantum speedups?”, Edited transcript of a rapporteur talk delivered at the 28th Solvay Physics Conference, Brussels, May 21, 2022 (2022).
 - [85] D. Hangleiter and J. Eisert, “Computational advantage of quantum random sampling”, Reviews of Modern Physics **95**, 035001.
 - [86] F. Arute et al., “Quantum supremacy using a programmable superconducting processor”, Nature **574**, 505–510 (2019).
 - [87] S. Aaronson and A. Arkhipov, Theory of Computing **9**, 143–252 (2013).
 - [88] M. J. Bremner, R. Jozsa, and D. J. Shepherd, “Classical simulation of commuting quantum computations implies collapse of the polynomial hierarchy”, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences **467**, 459–472 (2010).

- [89] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, R. Babbush, N. Ding, Z. Jiang, M. J. Bremner, J. M. Martinis, and H. Neven, “Characterizing quantum supremacy in near-term devices”, *Nature Physics* **14**, 595–600 (2018).
- [90] A. Bouland, B. Fefferman, C. Nirkhe, and U. Vazirani, “On the complexity and verification of quantum random circuit sampling”, *Nature Physics* **15**, 159–163 (2018).
- [91] H.-S. Zhong et al., “Quantum computational advantage using photons”, *Science* **370**, 1460–1463 (2020).
- [92] Y.-H. Deng et al., “Gaussian boson sampling with pseudo-photon-number-resolving detectors and quantum computational advantage”, *Physical Review Letters* **131**, 150601 (2023).
- [93] L. S. Madsen et al., “Quantum computational advantage with a programmable photonic processor”, *Nature* **606**, 75–81 (2022).
- [94] Y. Wu et al., “Strong quantum computational advantage using a superconducting quantum processor”, *Physical Review Letters* **127**, 180501 (2021).
- [95] Q. Zhu et al., “Quantum computational advantage via 60-qubit 24-cycle random circuit sampling”, *Science Bulletin* **67**, 240–245 (2022).
- [96] D. A. Abanin et al., “Observation of constructive interference at the edge of quantum ergodicity”, *Nature* **646**, 825–830 (2025).
- [97] S. R. White, “Density matrix formulation for quantum renormalization groups”, *Physical Review Letters* **69**, 2863–2866 (1992).
- [98] G. Vidal, “Efficient classical simulation of slightly entangled quantum computations”, *Physical Review Letters* **91**, 147902 (2003).
- [99] F. Verstraete, V. Murg, and J. Cirac, “Matrix product states, projected entangled pair states, and variational renormalization group methods for quantum spin systems”, *Advances in Physics* **57**, 143–224 (2008).
- [100] R. Orús, “A practical introduction to tensor networks: matrix product states and projected entangled pair states”, *Annals of Physics* **349**, 117–158 (2014).
- [101] S. Patra, S. S. Jahromi, S. Singh, and R. Orús, “Efficient tensor network simulation of ibm’s largest quantum processors”, *Physical Review Research* **6**, 013326 (2024).
- [102] T. Begušić, J. Gray, and G. K.-L. Chan, “Fast and converged classical simulations of evidence for the utility of quantum computing before fault tolerance”, *Science Advances* **10**, 10.1126/sciadv.adk4321 (2024).
- [103] D. Gottesman, “The Heisenberg representation of quantum computers”, *arXiv quant-ph/9807006* (1998).
- [104] S. Aaronson and D. Gottesman, “Improved simulation of stabilizer circuits”, *Physical Review A* **70**, 052328 (2004).
- [105] S. Bravyi and D. Gosset, “Improved classical simulation of quantum circuits dominated by Clifford gates”, *Physical Review Letters* **116**, 250501 (2016).
- [106] H. Pashayan, O. Reardon-Smith, K. Korzekwa, and S. D. Bartlett, “Fast estimation of outcome probabilities for quantum circuits”, *PRX Quantum* **3**, 020361 (2022).
- [107] D. Aharonov, X. Gao, Z. Landau, Y. Liu, and U. Vazirani, “A polynomial-time classical algorithm for noisy random circuit sampling”, in *Proceedings of the 55th annual acm symposium on theory of computing, STOC ’23* (2023), pp. 945–957.

- [108] G. González-García, J. I. Cirac, and R. Trivedi, “Pauli path simulations of noisy quantum circuits beyond average case”, *Quantum* **9**, 1730 (2025).
- [109] M. S. Rudolph, T. Jones, Y. Teng, A. Angrisani, and Z. Holmes, *Pauli propagation: a computational framework for simulating quantum systems*, 2025.
- [110] L. G. Valiant, “Quantum circuits that can be simulated classically in polynomial time”, *SIAM Journal on Computing* **31**, 1229–1254 (2002).
- [111] B. M. Terhal and D. P. DiVincenzo, “Classical simulation of noninteracting-fermion quantum circuits”, *Physical Review A* **65**, 032325 (2002).
- [112] E. Knill, *Fermionic linear optics and matchgates*, 2001.
- [113] J. Eisert, M. Cramer, and M. B. Plenio, “Colloquium: area laws for the entanglement entropy”, *Reviews of Modern Physics* **82**, 277–306 (2010).
- [114] M. Howard, J. Wallman, V. Veitch, and J. Emerson, “Contextuality supplies the ‘magic’ for quantum computation”, *Nature* **510**, 351–355 (2014).
- [115] V. Veitch, S. A. Hamed Mousavian, D. Gottesman, and J. Emerson, “The resource theory of stabilizer quantum computation”, *New Journal of Physics* **16**, 013009 (2014).
- [116] M. Hebenstreit, R. Jozsa, B. Kraus, S. Strelchuk, and M. Yoganathan, “All pure fermionic non-gaussian states are magic states for matchgate computations”, *Physical Review Letters* **123**, 080503 (2019).
- [117] L. Leone, S. F. E. Oliviero, and A. Hamma, “Stabilizer rényi entropy”, *Physical Review Letters* **128**, 050402 (2022).
- [118] A. Katabarwa, K. Gratsea, A. Caesura, and P. D. Johnson, “Early fault-tolerant quantum computing”, *PRX Quantum* **5**, 020101 (2024).
- [119] Z. Cai, R. Babbush, S. C. Benjamin, S. Endo, W. J. Huggins, Y. Li, J. R. McClean, and T. E. O’Brien, “Quantum error mitigation”, *Reviews of Modern Physics* **95**, 045005 (2023).
- [120] P. Hauke, F. M. Cucchietti, L. Tagliacozzo, I. Deutsch, and M. Lewenstein, “Can one trust quantum simulators?”, *Reports on Progress in Physics* **75**, 082401 (2012).
- [121] K. Temme, S. Bravyi, and J. M. Gambetta, “Error mitigation for short-depth quantum circuits”, *Physical Review Letters* **119**, 180509 (2017).
- [122] Y. Li and S. C. Benjamin, “Efficient variational quantum simulator incorporating active error minimization”, *Physical Review X* **7**, 021050 (2017).
- [123] E. van den Berg, Z. K. Mineev, A. Kandala, and K. Temme, “Probabilistic error cancellation with sparse Pauli–Lindblad models on noisy quantum processors”, *Nature Physics* **19**, 1116–1121 (2023).
- [124] W. J. Huggins, S. McArdle, T. E. O’Brien, J. Lee, N. C. Rubin, S. Boixo, K. B. Whaley, R. Babbush, and J. R. McClean, “Virtual distillation for quantum error mitigation”, *Physical Review X* **11**, 041036 (2021).
- [125] B. Koczor, “Exponential error suppression for near-term quantum devices”, *Physical Review X* **11**, 031057 (2021).
- [126] P. Czarnik, A. Arrasmith, P. J. Coles, and L. Cincio, “Error mitigation with Clifford quantum-circuit data”, *Quantum* **5**, 592 (2021).

- [127] X. Bonet-Monroig, R. Sagastizabal, M. Singh, and T. E. O’Brien, “Low-cost error mitigation by symmetry verification”, *Physical Review A* **98**, 062339 (2018).
- [128] E. Ballini, J. Mildenerberger, M. M. Wauters, and P. Hauke, “Symmetry verification for noisy quantum simulations of non-abelian lattice gauge theories”, *Quantum* **9**, 1802 (2025).
- [129] R. Takagi, S. Endo, S. Minagawa, and M. Gu, “Fundamental limits of quantum error mitigation”, *npj Quantum Information* **8**, 10.1038/s41534-022-00618-z (2022).
- [130] Y. Quek, D. Stilck França, S. Khatrri, J. J. Meyer, and J. Eisert, “Exponentially tighter bounds on limitations of quantum error mitigation”, *Nature Physics* **20**, 1648–1658 (2024).
- [131] E. L. Hahn, “Spin echoes”, *Physical Review* **80**, 580–594 (1950).
- [132] L. Viola and S. Lloyd, “Dynamical suppression of decoherence in two-state quantum systems”, *Physical Review A* **58**, 2733–2744 (1998).
- [133] L. Viola, E. Knill, and S. Lloyd, “Dynamical decoupling of open quantum systems”, *Physical Review Letters* **82**, 2417–2421 (1999).
- [134] K. Khodjasteh and D. A. Lidar, “Fault-tolerant quantum dynamical decoupling”, *Physical Review Letters* **95**, 180501 (2005).
- [135] A. M. Souza, G. A. Álvarez, and D. Suter, “Robust dynamical decoupling for quantum computing and quantum memory”, *Physical Review Letters* **106**, 240501 (2011).
- [136] G. T. Genov, D. Schraft, N. V. Vitanov, and T. Halfmann, “Arbitrarily accurate pulse sequences for robust dynamical decoupling”, *Physical Review Letters* **118**, 133202 (2017).
- [137] P. Wang, C.-Y. Luan, M. Qiao, M. Um, J. Zhang, Y. Wang, X. Yuan, M. Gu, J. Zhang, and K. Kim, “Single ion qubit with estimated coherence time exceeding one hour”, *Nature Communications* **12**, 10.1038/s41467-020-20330-w (2021).
- [138] X. Shi, J. Sinanan-Singh, K. DeBry, S. L. Todaro, I. L. Chuang, and J. Chiaverini, “Long-lived metastable-qubit memory”, *Physical Review A* **111**, 1020601 (2025).
- [139] N. Ezzell, B. Pokharel, L. Tewala, G. Quiroz, and D. A. Lidar, “Dynamical decoupling for superconducting qubits: a performance survey”, *Physical Review Applied* **20**, 064027 (2023).
- [140] V. Tripathi, H. Chen, M. Khezri, K.-W. Yip, E. Levenson-Falk, and D. A. Lidar, “Suppression of crosstalk in superconducting qubits using dynamical decoupling”, *Physical Review Applied* **18**, 024068 (2022).
- [141] N. Bar-Gill, L. Pham, A. Jarmola, D. Budker, and R. Walsworth, “Solid-state electronic spin coherence time approaching one second”, *Nature Communications* **4**, 10.1038/ncomms2771 (2013).
- [142] A. Jenkins, J. W. Lis, A. Senoo, W. F. McGrew, and A. M. Kaufman, “Ytterbium nuclear-spin qubits in an optical tweezer array”, *Physical Review X* **12**, 021027 (2022).
- [143] D. A. Lidar, “Towards fault tolerant adiabatic quantum computation”, *Physical Review Letters* **100**, 160506 (2008).
- [144] S. Nagies, C. Capecci, M. S. Benkner, J. Akram, S. Rubbert, D. Bantounas, M. Moeller, M. Johanning, and P. Hauke, “Practical noise mitigation for quantum annealing via dynamical decoupling: toward industry-relevant optimization using trapped ions”, *Phys. Rev. A* **113**, 012621 (2026).

-
- [145] R. Acharya et al., “Quantum error correction below the surface code threshold”, *Nature* **638**, 920–926 (2024).
 - [146] C. Ryan-Anderson, J. G. Bohnet, K. Lee, D. Gresh, A. Hankin, J. P. Gaebler, D. Francois, A. Chernoguzov, D. Lucchetti, N. C. Brown, T. M. Gatterman, S. K. Halit, K. Gilmore, J. A. Gerber, B. Neyenhuis, D. Hayes, and R. P. Stutz, “Realization of real-time fault-tolerant quantum error correction”, *Physical Review X* **11**, 041058 (2021).
 - [147] D. Bluvstein, S. J. Evered, A. A. Geim, S. H. Li, H. Zhou, T. Manovitz, S. Ebadi, M. Cain, M. Kalinowski, D. Hangleiter, J. P. Bonilla Ataides, N. Maskara, I. Cong, X. Gao, P. Sales Rodriguez, T. Karolyshyn, G. Semeghini, M. J. Gullans, M. Greiner, V. Vuletić, and M. D. Lukin, “Logical quantum processor based on reconfigurable atom arrays”, *Nature* **626**, 58–65 (2023).
 - [148] D. Bluvstein, A. A. Geim, S. H. Li, S. J. Evered, J. P. Bonilla Ataides, G. Baranes, A. Gu, T. Manovitz, M. Xu, M. Kalinowski, S. Majidy, C. Kokail, N. Maskara, E. C. Trapp, L. M. Stewart, S. Hollerith, H. Zhou, M. J. Gullans, S. F. Yelin, M. Greiner, V. Vuletić, M. Cain, and M. D. Lukin, “A fault-tolerant neutral-atom architecture for universal quantum computation”, *Nature* **649**, 39–46 (2025).
 - [149] S. Bravyi, A. W. Cross, J. M. Gambetta, D. Maslov, P. Rall, and T. J. Yoder, “High-threshold and low-overhead fault-tolerant quantum memory”, *Nature* **627**, 778–782 (2024).
 - [150] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, “Charge-insensitive qubit design derived from the cooper pair box”, *Physical Review A* **76**, 042319 (2007).
 - [151] V. E. Manucharyan, J. Koch, L. I. Glazman, and M. H. Devoret, “Fluxonium: single cooper-pair circuit free of charge offsets”, *Science* **326**, 113–116 (2009).
 - [152] F. Marxer et al., *Above 99.9% fidelity single-qubit gates, two-qubit gates, and readout in a single superconducting quantum device*, 2025.
 - [153] Y. Sung et al., “Realization of high-fidelity cz and zz -free iswap gates with a tunable coupler”, *Physical Review X* **11**, 021058 (2021).
 - [154] A. Kandala, K. X. Wei, S. Srinivasan, E. Magesan, S. Carnevale, G. A. Keefe, D. Klaus, O. Dial, and D. C. McKay, “Demonstration of a high-fidelity cnot gate for fixed-frequency transmons with engineered zz suppression”, *Physical Review Letters* **127**, 130501 (2021).
 - [155] L. Ding, M. Hays, Y. Sung, B. Kannan, J. An, A. Di Paolo, A. H. Karamlou, T. M. Hazard, K. Azar, D. K. Kim, B. M. Niedzielski, A. Melville, M. E. Schwartz, J. L. Yoder, T. P. Orlando, S. Gustavsson, J. A. Grover, K. Serniak, and W. D. Oliver, “High-fidelity, frequency-flexible two-qubit fluxonium gates with a transmon coupler”, *Physical Review X* **13**, 031035 (2023).
 - [156] J. Gambetta, *IBM Quantum System Two: the era of quantum utility is here*, IBM Quantum blog, <https://www.ibm.com/quantum/blog/quantum-roadmap-2033>, 2023.
 - [157] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, “Trapped-ion quantum computing: progress and challenges”, *Applied Physics Reviews* **6**, 10.1063/1.5088164 (2019).

- [158] J. I. Cirac and P. Zoller, “Quantum computations with cold trapped ions”, *Physical Review Letters* **74**, 4091–4094 (1995).
- [159] K. Mølmer and A. Sørensen, “Multiparticle entanglement of hot trapped ions”, *Physical Review Letters* **82**, 1835–1838 (1999).
- [160] A. Sørensen and K. Mølmer, “Quantum computation with ions in thermal motion”, *Physical Review Letters* **82**, 1971–1974 (1999).
- [161] F. Mintert and C. Wunderlich, “Ion-trap quantum logic using long-wavelength radiation”, *Physical Review Letters* **87**, 257904 (2001).
- [162] C. Wunderlich, “Conditional spin resonance with trapped ions”, in *Laser physics at the limits* (Springer Berlin Heidelberg, 2002), pp. 261–273.
- [163] C. Piltz, T. Sriarunothai, S. S. Ivanov, S. Wölk, and C. Wunderlich, “Versatile microwave-driven trapped ion spin system for quantum information processing”, *Science Advances* **2**, 10.1126/sciadv.1600093 (2016).
- [164] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas, “High-fidelity quantum logic gates using trapped-ion hyperfine qubits”, *Physical Review Letters* **117**, 060504 (2016).
- [165] J. P. Gaebler, T. R. Tan, Y. Lin, Y. Wan, R. Bowler, A. C. Keith, S. Glancy, K. Coakley, E. Knill, D. Leibfried, and D. J. Wineland, “High-fidelity universal gate set for $^9\text{Be}^+$ ion qubits”, *Physical Review Letters* **117**, 060505 (2016).
- [166] T. P. Harty, M. A. Sepiol, D. T. C. Allcock, C. J. Ballance, J. E. Tarlton, and D. M. Lucas, “High-fidelity trapped-ion quantum logic using near-field microwaves”, *Physical Review Letters* **117**, 140501 (2016).
- [167] 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, “High-fidelity laser-free universal control of trapped ion qubits”, *Nature* **597**, 209–213 (2021).
- [168] A. C. Hughes, R. Srinivas, C. M. Löschnauer, H. M. Knaack, R. Matt, C. J. Ballance, M. Malinowski, T. P. Harty, and R. T. Sutherland, *Trapped-ion two-qubit gates with >99.99% fidelity without ground-state cooling*, 2025.
- [169] M. Saffman, T. G. Walker, and K. Mølmer, “Quantum information with rydberg atoms”, *Reviews of Modern Physics* **82**, 2313–2363 (2010).
- [170] L. Henriot, L. Beguin, A. Signoles, T. Lahaye, A. Browaeys, G.-O. Reymond, and C. Jurczak, “Quantum computing with neutral atoms”, *Quantum* **4**, 327 (2020).
- [171] M. A. Norcia et al., “Iterative assembly of ^{171}Yb atom arrays with cavity-enhanced optical lattices”, *PRX Quantum* **5**, 030316 (2024).
- [172] B. W. Reichardt et al., *Fault-tolerant quantum computation with a neutral atom processor*, 2024.
- [173] E. Knill, R. Laflamme, and G. J. Milburn, “A scheme for efficient quantum computation with linear optics”, *Nature* **409**, 46–52 (2001).
- [174] S. Bartolucci, P. Birchall, H. Bombín, H. Cable, C. Dawson, M. Gimeno-Segovia, E. Johnston, K. Kieling, N. Nickerson, M. Pant, F. Pastawski, T. Rudolph, and C. Sparrow, “Fusion-based quantum computation”, *Nature Communications* **14**, 10.1038/s41467-023-36493-1 (2023).

-
- [175] K. Alexander et al., “A manufacturable platform for photonic quantum computing”, *Nature* **641**, 876–883 (2025).
 - [176] G. Burkard, T. D. Ladd, A. Pan, J. M. Nichol, and J. R. Petta, “Semiconductor spin qubits”, *Reviews of Modern Physics* **95**, 025003 (2023).
 - [177] A. Noiri, K. Takeda, T. Nakajima, T. Kobayashi, A. Sammak, G. Scappucci, and S. Tarucha, “Fast universal quantum gate above the fault-tolerance threshold in silicon”, *Nature* **601**, 338–342 (2022).
 - [178] X. Xue, M. Russ, N. Samkharadze, B. Undseth, A. Sammak, G. Scappucci, and L. M. K. Vandersypen, “Quantum logic with spin qubits crossing the surface code threshold”, *Nature* **601**, 343–347 (2022).
 - [179] M. W. Doherty, N. B. Manson, P. Delaney, F. Jelezko, J. Wrachtrup, and L. C. Hollenberg, “The nitrogen-vacancy colour centre in diamond”, *Physics Reports* **528**, 1–45 (2013).
 - [180] M. Aghaee et al., “Interferometric single-shot parity measurement in InAs–Al hybrid devices”, *Nature* **638**, 651–655 (2025).
 - [181] H. F. Legg, *Comment on “Interferometric single-shot parity measurement in InAs–Al hybrid devices”*, *Microsoft quantum, Nature 638, 651–655 (2025)*, 2025.
 - [182] T. Kadowaki and H. Nishimori, “Quantum annealing in the transverse ising model”, *Phys. Rev. E* **58**, 5355–5363 (1998).
 - [183] T. Albash and D. A. Lidar, “Adiabatic quantum computation”, *Reviews of Modern Physics* **90**, 015002 (2018).
 - [184] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, “Optimization by simulated annealing”, *Science* **220**, 671–680 (1983).
 - [185] S. Jansen, M.-B. Ruskai, and R. Seiler, “Bounds for the adiabatic approximation with applications to quantum computation”, *Journal of Mathematical Physics* **48**, 102111 (2007).
 - [186] D. A. Lidar, A. T. Rezakhani, and A. Hamma, “Adiabatic approximation with exponential accuracy for many-body systems and quantum computation”, *Journal of Mathematical Physics* **50**, 10.1063/1.3236685 (2009).
 - [187] M. H. S. Amin, “Consistency of the adiabatic theorem”, *Physical Review Letters* **102**, 220401 (2009).
 - [188] D. Cheung, P. Høyer, and N. Wiebe, “Improved error bounds for the adiabatic approximation”, *Journal of Physics A: Mathematical and Theoretical* **44**, 415302 (2011).
 - [189] A. Lucas, “Ising formulations of many NP problems”, *Frontiers in Physics* **2**, 10.3389/fphy.2014.00005 (2014).
 - [190] S. Bravyi, D. P. DiVincenzo, R. I. Oliveira, and B. M. Terhal, “The complexity of stoquastic local hamiltonian problems”, 10.48550/ARXIV.QUANT-PH/0606140 (2006).
 - [191] J. D. Biamonte and P. J. Love, “Realizable Hamiltonians for universal adiabatic quantum computers”, *Phys. Rev. A* **78**, 012352 (2008).
 - [192] T. Jörg, F. Krzakala, J. Kurchan, A. C. Maggs, and J. Pujos, “Energy gaps in quantum first-order mean-field-like transitions: the problems that quantum annealing cannot solve”, *EPL (Europhysics Letters)* **89**, 40004 (2010).

- [193] A. P. Young, S. Knysh, and V. N. Smelyanskiy, “First-order phase transition in the quantum adiabatic algorithm”, *Physical Review Letters* **104**, 020502 (2010).
- [194] N. Chancellor, S. Zohren, and P. A. Warburton, “Circuit design for multi-body interactions in superconducting quantum annealing systems with applications to a scalable architecture”, *npj Quantum Information* **3**, 10.1038/s41534-017-0022-6 (2017).
- [195] Y. Seki and H. Nishimori, “Quantum annealing with antiferromagnetic fluctuations”, *Physical Review E* **85**, 051112 (2012).
- [196] H. Nishimori and K. Takada, “Exponential enhancement of the efficiency of quantum annealing by non-stoquastic hamiltonians”, *Frontiers in ICT* **4**, 10.3389/fict.2017.00002 (2017).
- [197] L. Hormozi, E. W. Brown, G. Carleo, and M. Troyer, “Nonstoquastic Hamiltonians and quantum annealing of an Ising spin glass”, *Physical Review B* **95**, 184416 (2017).
- [198] A. D. King et al., “Quantum critical dynamics in a 5,000-qubit programmable spin glass”, *Nature* **617**, 61–66 (2023).
- [199] S. Ebadi et al., “Quantum optimization of maximum independent set using rydberg atom arrays”, *Science* **376**, 1209–1215 (2022).
- [200] R. Barends, A. Shabani, L. Lamata, J. Kelly, A. Mezzacapo, U. L. Heras, R. Babbush, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, E. Lucero, A. Megrant, J. Y. Mutus, M. Neeley, C. Neill, P. J. J. O’Malley, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, E. Solano, H. Neven, and J. M. Martinis, “Digitized adiabatic quantum computing with a superconducting circuit”, *Nature* **534**, 222–226 (2016).
- [201] J. W. Britton, B. C. Sawyer, A. C. Keith, C.-C. J. Wang, J. K. Freericks, H. Uys, M. J. Biercuk, and J. J. Bollinger, “Engineered two-dimensional ising interactions in a trapped-ion quantum simulator with hundreds of spins”, *Nature* **484**, 489–492 (2012).
- [202] C. Monroe, W. C. Campbell, L.-M. Duan, Z.-X. Gong, A. V. Gorshkov, P. W. Hess, R. Islam, K. Kim, N. M. Linke, G. Pagano, P. Richerme, C. Senko, and N. Y. Yao, “Programmable quantum simulations of spin systems with trapped ions”, *Reviews of Modern Physics* **93**, 025001 (2021).
- [203] S. Nagies, K. T. Geier, J. Akram, D. Bantounas, M. Johanning, and P. Hauke, “Boosting quantum annealing performance through direct polynomial unconstrained binary optimization”, *Quantum Science and Technology* **10**, 035008 (2025).
- [204] C. Capecchi, S. Nagies, N. D. Varikuti, and P. Hauke, *Quantum resources in non-stoquastic quantum annealing*, 2026.
- [205] S. Nagies, K. T. Geier, J. Akram, J. Okamoto, D. Bantounas, C. Wunderlich, M. Johanning, and P. Hauke, “The role of higher-order terms in trapped-ion quantum computing with magnetic gradient induced coupling”, *Quantum Science and Technology* **10**, 025051 (2025).
- [206] L. S. Brown and G. Gabrielse, “Geonium theory: physics of a single electron or ion in a Penning trap”, *Rev. Mod. Phys.* **58**, 233–311 (1986).
- [207] H. G. Dehmelt, “Experiments with an isolated subatomic particle at rest”, *Rev. Mod. Phys.* **62**, Nobel Lecture, 525–530 (1990).
- [208] A. M. Polloreno, A. M. Rey, and J. J. Bollinger, “Individual qubit addressing of rotating ion crystals in a Penning trap”, *Phys. Rev. Research* **4**, 033076 (2022).

- [209] M. J. Biercuk, H. Uys, A. P. VanDevender, N. Shiga, W. M. Itano, and J. J. Bollinger, “High-fidelity quantum control using ion crystals in a Penning trap”, *Quantum Inf. Comput.* **9**, 920–949 (2009).
- [210] J. G. Bohnet, B. C. Sawyer, J. W. Britton, M. L. Wall, A. M. Rey, M. Foss-Feig, and J. J. Bollinger, “Quantum spin dynamics and entanglement generation with hundreds of trapped ions”, *Science* **352**, 1297–1301 (2016).
- [211] M. Gärttner, J. G. Bohnet, A. Safavi-Naini, M. L. Wall, J. J. Bollinger, and A. M. Rey, “Measuring out-of-time-order correlations and multiple quantum spectra in a trapped-ion quantum magnet”, *Nat. Phys.* **13**, 781–786 (2017).
- [212] S. Jain, T. Sägger, P. Hrmo, C. Torkzaban, M. Stadler, R. Oswald, C. Axline, A. Bautista-Salvador, C. Ospelkaus, D. Kienzler, and J. Home, “Penning micro-trap for quantum computing”, *Nature* **627**, 510–514 (2024).
- [213] S. Jain, J. Alonso, M. Grau, and J. P. Home, “Scalable arrays of micro-Penning traps for quantum computing and simulation”, *Phys. Rev. X* **10**, 031027 (2020).
- [214] W. Paul, “Electromagnetic traps for charged and neutral particles”, *Reviews of Modern Physics* **62**, 531–540 (1990).
- [215] D. J. Berkeland, J. D. Miller, J. C. Bergquist, W. M. Itano, and D. J. Wineland, “Minimization of ion micromotion in a paul trap”, *Journal of Applied Physics* **83**, 5025–5033 (1998).
- [216] D. Leibfried, B. DeMarco, V. Meyer, D. Lucas, M. Barrett, J. Britton, W. M. Itano, B. Jelenković, C. Langer, T. Rosenband, and D. J. Wineland, “Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate”, *Nature* **422**, 412–415 (2003).
- [217] H. Häffner, C. Roos, and R. Blatt, “Quantum computing with trapped ions”, *Physics Reports* **469**, 155–203 (2008).
- [218] T. P. Harty, D. T. C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas, “High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit”, *Physical Review Letters* **113**, 220501 (2014).
- [219] D. James, “Quantum dynamics of cold trapped ions with application to quantum computation”, *Applied Physics B: Lasers and Optics* **66**, 181–190 (1998).
- [220] J. Chiaverini, R. B. Blakestad, J. Britton, J. D. Jost, C. Langer, D. Leibfried, R. Ozeri, and D. J. Wineland, “Surface-electrode architecture for ion-trap quantum information processing”, (2005).
- [221] S. Seidelin, J. Chiaverini, R. Reichle, J. Bollinger, D. Leibfried, J. Britton, J. Wesenberg, R. Blakestad, R. Epstein, D. Hume, W. Itano, J. Jost, C. Langer, R. Ozeri, N. Shiga, and D. Wineland, “Microfabricated surface-electrode ion trap for scalable quantum information processing”, *Physical Review Letters* **96**, 253003 (2006).
- [222] A. Ransford et al., *Helios: a 98-qubit trapped-ion quantum computer*, 2025.
- [223] S. Zippilli, M. Johanning, S. M. Giampaolo, C. Wunderlich, and F. Illuminati, “Adiabatic quantum simulation with a segmented ion trap: application to long-distance entanglement in quantum spin systems”, *Physical Review A* **89**, 042308 (2014).
- [224] R. Bowler, J. Gaebler, Y. Lin, T. R. Tan, D. Hanneke, J. D. Jost, J. P. Home, D. Leibfried, and D. J. Wineland, “Coherent diabatic ion transport and separation in a multizone trap array”, *Physical Review Letters* **109**, 080502 (2012).

- [225] D. Kielpinski, C. Monroe, and D. J. Wineland, “Architecture for a large-scale ion-trap quantum computer”, *Nature* **417**, 709–711 (2002).
- [226] Z. D. Romaszko, S. Hong, M. Siegele, R. K. Puddy, F. R. Lebrun-Gallagher, S. Weidt, and W. K. Hensinger, “Engineering of microfabricated ion traps and integration of advanced on-chip features”, *Nature Reviews Physics* **2**, 285–299 (2020).
- [227] C. Ospelkaus, U. Warring, Y. Colombe, K. R. Brown, J. M. Amini, D. Leibfried, and D. J. Wineland, “Microwave quantum logic gates for trapped ions”, *Nature* **476**, 181–184 (2011).
- [228] M. C. Smith, A. D. Leu, K. Miyanishi, M. F. Gely, and D. M. Lucas, “Single-qubit gates with errors at the 10^{-7} level”, *Physical Review Letters* **134**, 10.1103/42w2-6ccy (2025).
- [229] C. Marquet, F. Schmidt-Kaler, and D. James, “Phonon-phonon interactions due to non-linear effects in a linear ion trap”, *Applied Physics B: Lasers and Optics* **76**, 199–208 (2003).
- [230] J. P. Schiffer, “Phase transitions in anisotropically confined ionic crystals”, *Physical Review Letters* **70**, 818–821 (1993).
- [231] D. G. Enzer, M. M. Schauer, J. J. Gomez, M. S. Gulley, M. H. Holzscheiter, P. G. Kwiatt, S. K. Lamoreaux, C. G. Peterson, V. D. Sandberg, D. Tupa, A. G. White, R. J. Hughes, and D. F. V. James, “Observation of power-law scaling for phase transitions in linear trapped ion crystals”, *Physical Review Letters* **85**, 2466 (2000).
- [232] B. Merkel, K. Thirumalai, J. E. Tarlton, V. M. Schäfer, C. J. Ballance, T. P. Harty, and D. M. Lucas, “Magnetic field stabilization system for atomic physics experiments”, *Review of Scientific Instruments* **90**, 10.1063/1.5080093 (2019).
- [233] M. Johanning, A. F. Varón, and C. Wunderlich, “Quantum simulations with cold trapped ions”, *Journal of Physics B: Atomic, Molecular and Optical Physics* **42**, 154009 (2009).
- [234] A. Khromova, C. Piltz, B. Scharfenberger, T. F. Gloger, M. Johanning, A. F. Varón, and C. Wunderlich, “Designer spin pseudomolecule implemented with trapped ions in a magnetic gradient”, *Physical Review Letters* **108**, 220502 (2012).
- [235] P. Baßler, M. Zipper, C. Cedzich, M. Heinrich, P. H. Huber, M. Johanning, and M. Kliesch, “Synthesis of and compilation with time-optimal multi-qubit gates”, *Quantum* **7**, 984 (2023).
- [236] P. Huber, J. Haber, P. Barthel, J. J. García-Ripoll, E. Torrontegui, and C. Wunderlich, *Realization of a quantum perceptron gate with trapped ions*, 2021.
- [237] C. Balzer, A. Braun, T. Hannemann, C. Paape, M. Ettler, W. Neuhauser, and C. Wunderlich, “Electrodynamically trapped Yb^+ ions for quantum information processing”, *Physical Review A* **73**, 041407 (2006).
- [238] S. Weidt, J. Randall, S. C. Webster, K. Lake, A. E. Webb, I. Cohen, T. Navickas, B. Lekitsch, A. Retzker, and W. K. Hensinger, “Trapped-ion quantum logic with global radiation fields”, *Physical Review Letters* **117**, 220501 (2016).
- [239] C. Piltz, B. Scharfenberger, A. Khromova, A. F. Varón, and C. Wunderlich, “Protecting conditional quantum gates by robust dynamical decoupling”, *Physical Review Letters* **110**, 200501 (2013).

-
- [240] N. Timoney, I. Baumgart, M. Johanning, A. F. Varón, M. B. Plenio, A. Retzker, and C. Wunderlich, “Quantum gates and memory using microwave-dressed states”, *Nature* **476**, 185–188 (2011).
 - [241] T. Ruster, C. T. Schmiegelow, H. Kaufmann, C. Warschburger, F. Schmidt-Kaler, and U. G. Poschinger, “A long-lived Zeeman trapped-ion qubit”, *Applied Physics B* **122**, 10.1007/s00340-016-6527-4 (2016).
 - [242] G. Breit and I. I. Rabi, “Measurement of nuclear spin”, *Physical Review* **38**, 2082–2083 (1931).
 - [243] D. Xu and J. Cao, “Non-canonical distribution and non-equilibrium transport beyond weak system-bath coupling regime: a polaron transformation approach”, *Frontiers of Physics* **11**, 10.1007/s11467-016-0540-2 (2016).
 - [244] D. Porras and J. I. Cirac, “Effective quantum spin systems with trapped ions”, *Physical Review Letters* **92**, 207901 (2004).
 - [245] P. Jurcevic, B. P. Lanyon, P. Hauke, C. Hempel, P. Zoller, R. Blatt, and C. F. Roos, “Quasiparticle engineering and entanglement propagation in a quantum many-body system”, *Nature* **511**, 202–205 (2014).
 - [246] J. Zhang, G. Pagano, P. W. Hess, A. Kyprianidis, P. Becker, H. Kaplan, A. V. Gorshkov, Z.-X. Gong, and C. Monroe, “Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator”, *Nature* **551**, 601–604 (2017).
 - [247] M. Siegele-Brown, S. Hong, F. R. Lebrun-Gallagher, S. J. Hile, S. Weidt, and W. K. Hensinger, “Fabrication of surface ion traps with integrated current carrying wires enabling high magnetic field gradients”, *Quantum Science and Technology* **7**, 034003 (2022).
 - [248] A. Sørensen and K. Mølmer, “Entanglement and quantum computation with ions in thermal motion”, *Physical Review A* **62**, 022311 (2000).
 - [249] P. Richerme, Z.-X. Gong, A. Lee, C. Senko, J. Smith, M. Foss-Feig, S. Michalakis, A. V. Gorshkov, and C. Monroe, “Non-local propagation of correlations in quantum systems with long-range interactions”, *Nature* **511**, 198–201 (2014).
 - [250] P. Nevado and D. Porras, “Hidden frustrated interactions and quantum annealing in trapped-ion spin-phonon chains”, *Physical Review A* **93**, 013625 (2016).
 - [251] N. Trautmann and P. Hauke, “Trapped-ion quantum simulation of excitation transport: disordered, noisy, and long-range connected quantum networks”, *Physical Review A* **97**, 023606 (2018).
 - [252] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information* (Cambridge University Press, Cambridge, 2000).
 - [253] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, “Efficient Z gates for quantum computing”, *Physical Review A* **96**, 022330 (2017).
 - [254] F. Schmidt-Kaler, H. Häffner, M. Riebe, S. Gulde, G. P. T. Lancaster, T. Deuschle, C. Becher, C. F. Roos, J. Eschner, and R. Blatt, “Realization of the Cirac–Zoller controlled-NOT quantum gate”, *Nature* **422**, 408–411 (2003).
 - [255] P. J. Lee, K.-A. Brickman, L. Deslauriers, P. C. Haljan, L.-M. Duan, and C. Monroe, “Phase control of trapped ion quantum gates”, *Journal of Optics B: Quantum and Semiclassical Optics* **7**, S371–S383 (2005).

- [256] C. Monroe, *Primer on Mølmer-Sørensen gates in trapped ions*, Lecture notes, Boulder Summer School for Condensed Matter and Materials Physics, 2020.
- [257] S.-L. Zhu, C. Monroe, and L.-M. Duan, “Trapped ion quantum computation with transverse phonon modes”, *Phys. Rev. Lett.* **97**, 050505 (2006).
- [258] R. Ozeri, W. M. Itano, R. B. Blakestad, J. Britton, J. Chiaverini, B. DeMarco, D. B. Hume, D. Leibfried, C. Langer, R. Reichle, S. Seidelin, J. H. Wesenberg, and D. J. Wineland, “Errors in trapped-ion quantum gates due to spontaneous photon scattering”, *Physical Review A* **75**, 042329 (2007).
- [259] B. Lekitsch, S. Weidt, A. G. Fowler, K. Mølmer, S. J. Devitt, C. Wunderlich, and W. K. Hensinger, “Blueprint for a microwave trapped ion quantum computer”, *Science Advances* **3**, 10.1126/sciadv.1601540 (2017).
- [260] R. Srinivas, S. C. Burd, R. T. Sutherland, A. C. Wilson, D. J. Wineland, D. Leibfried, D. T. C. Allcock, and D. H. Slichter, “Trapped-ion spin-motion coupling with microwaves and a near-motional oscillating magnetic field gradient”, *Physical Review Letters* **122**, 163201 (2019).
- [261] M. Loewen and C. Wunderlich, “Thermisch angeregte Ionen-, Moleküle“, *Verhandl. DPG 2004 (VI)* 39, 7/87”, Article (2004).
- [262] I. Pogorelov, T. Feldker, C. D. Marciniak, L. Postler, G. Jacob, O. Kriegelsteiner, V. Podlesnic, M. Meth, V. Negnevitsky, M. Stadler, B. Höfer, C. Wächter, K. Lakhmanskiy, R. Blatt, P. Schindler, and T. Monz, “Compact ion-trap quantum computing demonstrator”, *PRX Quantum* **2**, 020343 (2021).
- [263] P. Schindler, D. Nigg, T. Monz, J. T. Barreiro, E. Martinez, S. X. Wang, S. Quint, M. F. Brandl, V. Nebendahl, C. F. Roos, M. Chwalla, M. Hennrich, and R. Blatt, “A quantum information processor with trapped ions”, *New Journal of Physics* **15**, 123012 (2013).
- [264] H. Rohde, S. T. Gulde, C. F. Roos, P. A. Barton, D. Leibfried, J. Eschner, F. Schmidt-Kaler, and R. Blatt, “Sympathetic ground-state cooling and coherent manipulation with two-ion crystals”, *Journal of Optics B: Quantum and Semiclassical Optics* **3**, S34–S41 (2001).
- [265] T. Sriarunothai, G. S. Giri, S. Wölk, and C. Wunderlich, “Radio frequency sideband cooling and sympathetic cooling of trapped ions in a static magnetic field gradient”, *Journal of Modern Optics* **65**, 560–567 (2017).
- [266] K. Sosnova, A. Carter, and C. Monroe, “Character of motional modes for entanglement and sympathetic cooling of mixed-species trapped-ion chains”, *Physical Review A* **103**, 012610 (2021).
- [267] R. Schmied, T. Roscilde, V. Murg, D. Porras, and J. I. Cirac, “Quantum phases of trapped ions in an optical lattice”, *New Journal of Physics* **10**, 045017 (2008).
- [268] M. Johanning, “Isospaced linear ion strings”, *Applied Physics B* **122**, 10.1007/s00340-016-6340-0 (2016).
- [269] A. Bermudez, D. Porras, and M. A. Martin-Delgado, “Competing many-body interactions in systems of trapped ions”, *Physical Review A* **79**, 060303 (2009).
- [270] P. Blümler and H. Soltner, “Practical concepts for design, construction and application of halbach magnets in magnetic resonance”, *Applied Magnetic Resonance* **54**, 1701–1739 (2023).

-
- [271] D. González-Cuadra, P. R. Grzybowski, A. Dauphin, and M. Lewenstein, “Strongly correlated bosons on a dynamical lattice”, *Physical Review Letters* **121**, 090402 (2018).
 - [272] D. González-Cuadra, A. Dauphin, P. R. Grzybowski, P. Wójcik, M. Lewenstein, and A. Bermudez, “Symmetry-breaking topological insulators in the $\mathbb{Z}_2$ Bose-Hubbard model”, *Physical Review B* **99**, 045139 (2019).
 - [273] D. Yang, G. S. Giri, M. Johanning, C. Wunderlich, P. Zoller, and P. Hauke, “Analog quantum simulation of (1+1)-dimensional lattice QED with trapped ions”, *Physical Review A* **94**, 052321 (2016).
 - [274] B. Andrade, Z. Davoudi, T. Graß, M. Hafezi, G. Pagano, and A. Seif, “Engineering an effective three-spin hamiltonian in trapped-ion systems for applications in quantum simulation”, *Quantum Science and Technology* **7**, 034001 (2022).
 - [275] J. Mildenerger, W. Mruczkiewicz, J. C. Halimeh, Z. Jiang, and P. Hauke, “Confinement in a $\mathbb{Z}_2$ lattice gauge theory on a quantum computer”, *Nature Physics* **21**, 312–317 (2025).
 - [276] W. Lechner, P. Hauke, and P. Zoller, “A quantum annealing architecture with all-to-all connectivity from local interactions”, *Science Advances* **1**, 10.1126/sciadv.1500838 (2015).
 - [277] P. Hauke, D. Marcos, M. Dalmonte, and P. Zoller, “Quantum simulation of a lattice schwinger model in a chain of trapped ions”, *Physical Review X* **3**, 041018 (2013).
 - [278] K. S. C. Decker, C. Karrasch, J. Eisert, and D. M. Kennes, “Floquet engineering topological many-body localized systems”, *Physical Review Letters* **124**, 190601 (2020).
 - [279] O. Katz, M. Cetina, and C. Monroe, “N-body interactions between trapped ion qubits via spin-dependent squeezing”, *Physical Review Letters* **129**, 063603 (2022).
 - [280] J. Roland and N. J. Cerf, “Quantum search by local adiabatic evolution”, *Physical Review A* **65**, 042308 (2002).
 - [281] A. T. Rezakhani, W.-J. Kuo, A. Hamma, D. A. Lidar, and P. Zanardi, “Quantum adiabatic brachistochrone”, *Physical Review Letters* **103**, 080502 (2009).
 - [282] A. Perdomo-Ortiz, S. E. Venegas-Andraca, and A. Aspuru-Guzik, “A study of heuristic guesses for adiabatic quantum computation”, *Quantum Information Processing* **10**, 33–52 (2010).
 - [283] N. Chancellor, “Modernizing quantum annealing using local searches”, *New Journal of Physics* **19**, 023024 (2017).
 - [284] M. Ohkuwa, H. Nishimori, and D. A. Lidar, “Reverse annealing for the fully connected p-spin model”, *Physical Review A* **98**, 022314 (2018).
 - [285] Y. Susa, Y. Yamashiro, M. Yamamoto, and H. Nishimori, “Exponential speedup of quantum annealing by inhomogeneous driving of the transverse field”, *Journal of the Physical Society of Japan* **87**, 023002 (2018).
 - [286] J. I. Adame and P. L. McMahon, “Inhomogeneous driving in quantum annealers can result in orders-of-magnitude improvements in performance”, *Quantum Science and Technology* **5**, 035011 (2020).
 - [287] A. del Campo, “Shortcuts to adiabaticity by counterdiabatic driving”, *Physical Review Letters* **111**, 100502 (2013).

- [288] D. Sels and A. Polkovnikov, “Minimizing irreversible losses in quantum systems by local counterdiabatic driving”, *Proceedings of the National Academy of Sciences* **114**, 10.1073/pnas.1619826114 (2017).
- [289] D. Guéry-Odelin, A. Ruschhaupt, A. Kiely, E. Torrontegui, S. Martínez-Garaot, and J. G. Muga, “Shortcuts to adiabaticity: concepts, methods, and applications”, *Reviews of Modern Physics* **91**, 045001 (2019).
- [290] I. Čepaitė, A. Polkovnikov, A. J. Daley, and C. W. Duncan, “Counterdiabatic optimized local driving”, *PRX Quantum* **4**, 010312 (2023).
- [291] L. P. García-Pintos, M. Sahasrabudhe, and C. Arenz, “Tighter lower bounds on quantum annealing times”, (2024).
- [292] A. Bottarelli, M. G. de Andoin, P. Chandarana, K. Paul, X. Chen, M. Sanz, and P. Hauke, “Symmetry-enhanced counterdiabatic quantum algorithm for qudits”, (2024).
- [293] B. Altshuler, H. Krovi, and J. Roland, “Anderson localization makes adiabatic quantum optimization fail”, *Proceedings of the National Academy of Sciences* **107**, 12446–12450 (2010).
- [294] T. Albash, “Role of nonstoquastic catalysts in quantum adiabatic optimization”, *Physical Review A* **99**, 042334 (2019).
- [295] H. Lipkin, N. Meshkov, and A. Glick, “Validity of many-body approximation methods for a solvable model”, *Nuclear Physics* **62**, 188–198 (1965).
- [296] N. Meshkov, A. Glick, and H. Lipkin, “Validity of many-body approximation methods for a solvable model”, *Nuclear Physics* **62**, 199–210 (1965).
- [297] A. Glick, H. Lipkin, and N. Meshkov, “Validity of many-body approximation methods for a solvable model”, *Nuclear Physics* **62**, 211–224 (1965).
- [298] V. Bapst and G. Semerjian, “On quantum mean-field models and their quantum annealing”, *Journal of Statistical Mechanics: Theory and Experiment* **2012**, P06007 (2012).
- [299] W. Luo, J. Xing, A. Milan, X. Zhang, W. Liu, and T.-K. Kim, “Multiple object tracking: a literature review”, *Artificial Intelligence* **293**, 103448 (2021).
- [300] Z.-Q. Zhao, P. Zheng, S.-t. Xu, and X. Wu, “Object detection with deep learning: a review”, *IEEE transactions on neural networks and learning systems* **30**, 3212–3232 (2019).
- [301] R. Galian, T. Hamm, and S. Ordyniak, “The complexity of object association in multiple object tracking”, in *Proceedings of the aaai conference on artificial intelligence*, Vol. 35, 2 (2021), pp. 1388–1396.
- [302] J.-N. Zaech, A. Liniger, M. Danelljan, D. Dai, and L. Van Gool, “Adiabatic quantum computing for multi object tracking”, in *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2022).
- [303] Y. Ihara, “Enhancing multiple object tracking accuracy via quantum annealing”, *Scientific Reports* **15**, 24294 (2025).
- [304] M. Andriluka, S. Roth, and B. Schiele, “Monocular 3d pose estimation and tracking by detection”, in *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (Ieee, 2010)*, pp. 623–630.
- [305] L. Leal-Taixé, A. Milan, I. Reid, S. Roth, and K. Schindler, *Motchallenge 2015: towards a benchmark for multi-target tracking*, 2015.

-
- [306] Z. Zheng, X. Yang, Z. Yu, L. Zheng, Y. Yang, and J. Kautz, “Joint discriminative and generative learning for person re-identification”, in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition* (2019), pp. 2138–2147.
 - [307] A. Hornakova, T. Kaiser, P. Swoboda, M. Rolinek, B. Rosenhahn, and R. Henschel, “Making higher order mot scalable: an efficient approximate solver for lifted disjoint paths”, in *Proceedings of the IEEE/CVF international conference on computer vision* (2021), pp. 6330–6340.
 - [308] M. Seelbach Benkner, V. Golyanik, C. Theobalt, and M. Moeller, “Adiabatic quantum graph matching with permutation matrix constraints”, in *2020 international conference on 3d vision (3dv)* (IEEE, 2020), pp. 583–592.
 - [309] E. Alessandroni, S. Ramos-Calderer, I. Roth, E. Traversi, and L. Aolita, “Alleviating the quantum big- M problem”, (2023).
 - [310] *MOT Challenge - Results* — *motchallenge.net*, <https://motchallenge.net/results/MOT20Det/>, [Accessed 15-08-2025].
 - [311] L. V. Kantorovich, “Mathematical methods of organizing and planning production”, *Management science* **6**, 366–422 (1960).
 - [312] T. Bontekoe, F. Phillipson, and W. v. d. Schoot, “Translating constraints into qubos for the quadratic knapsack problem”, in *International conference on computational science* (Springer, 2023), pp. 90–107.
 - [313] K. Ohno, T. Shirai, and N. Togawa, “Toward practical benchmarks of Ising machines: a case study on the quadratic knapsack problem”, *IEEE Access* **12**, 97678–97690 (2024).
 - [314] E. Güney, J. Ehrental, and T. Hanne, “QUBO Formulations and Characterization of Penalty Parameters for the Multi-Knapsack Problem”, *IEEE Access* **13**, 47086–47098 (2025).
 - [315] R. A. Quintero and L. F. Zuluaga, “Characterizing and benchmarking QUBO reformulations of the knapsack problem”, Technical Report. Department of Industrial and Systems Engineering, Lehigh University, Tech. Rep. (2021).
 - [316] A. Montanez-Barrera, D. Willsch, A. Maldonado-Romo, and K. Michielsen, “Unbalanced penalization: a new approach to encode inequality constraints of combinatorial problems for quantum optimization algorithms”, *Quantum Science and Technology*, Volume 9, Number 2 (2024), 10.1088/2058-9565/ad35e4 (2022).
 - [317] A. Bottarelli, S. Schmitt, and P. Hauke, “Inequality constraints in variational quantum circuits with qudits”, *Physical Review Research* **7**, 10.1103/3196-41xf (2025).
 - [318] F. Glover, G. Kochenberger, and Y. Du, *A tutorial on formulating and using QUBO models*, 2019.
 - [319] B. Aspvall, M. F. Plass, and R. E. Tarjan, “A linear-time algorithm for testing the truth of certain quantified boolean formulas”, *Information Processing Letters* **14**, 195 (1982).
 - [320] S. A. Cook, “The complexity of theorem-proving procedures”, in *Proceedings of the third annual acm symposium on theory of computing - stoc '71*, STOC '71 (1971).
 - [321] R. M. Karp, “Reducibility among combinatorial problems”, in *Complexity of computer computations* (Springer US, 1972), pp. 85–103.

- [322] S. Aaronson, “Bqp and the polynomial hierarchy”, in Proceedings of the forty-second acm symposium on theory of computing, STOC’10 (June 2010).
- [323] B. Heim, T. F. Rønnow, S. V. Isakov, and M. Troyer, “Quantum versus classical annealing of ising spin glasses”, *Science* **348**, 215–217 (2015).
- [324] D. Willsch, M. Willsch, C. D. Gonzalez Calaza, F. Jin, H. De Raedt, M. Svensson, and K. Michielsen, “Benchmarking advantage and d-wave 2000q quantum annealers with exact cover problems”, *Quantum Information Processing* **21**, 10.1007/s11128-022-03476-y (2022).
- [325] M. Bernaschi, I. González-Adalid Pemartín, V. Martín-Mayor, and G. Parisi, “The quantum transition of the two-dimensional ising spin glass”, *Nature* **631**, 749–754 (2024).
- [326] M. Mézard, G. Parisi, and R. Zecchina, “Analytic and algorithmic solution of random satisfiability problems”, *Science* **297**, 812–815 (2002).
- [327] M. Žnidarič, “Scaling of the running time of the quantum adiabatic algorithm for propositional satisfiability”, *Physical Review A* **71**, 062305 (2005).
- [328] C.-F. C. Chang-Yu Hsieh Robert Amador, “Exploration of hard to solve 3-sat problems”, *INFORMATION TECHNOLOGY IN INDUSTRY* **7**, 10.17762/itii.v7i2.67 (2021).
- [329] D-Wave Systems, <https://www.dwavesys.com/>.
- [330] A. Verma, M. Lewis, and G. Kochenberger, *Efficient QUBO transformation for higher degree pseudo boolean functions*, 2021.
- [331] V. Choi, *Adiabatic quantum algorithms for the NP-complete maximum-weight independent set, exact cover and 3SAT problems*, 2010.
- [332] R. Herrman, L. Treffert, J. Ostrowski, P. C. Lotshaw, T. S. Humble, and G. Siopsis, “Globally optimizing qaoa circuit depth for constrained optimization problems”, *Algorithms* **14**, 294 (2021).
- [333] M. K. Bashar and N. Shukla, “Designing ising machines with higher order spin interactions and their application in solving combinatorial optimization”, *Scientific Reports* **13**, 10.1038/s41598-023-36531-4 (2023).
- [334] D. Y. Kang, T. Kelly, D. Kühn, A. Methuku, and D. Osthus, “Graph and hypergraph colouring via nibble methods: a survey”, in *European congress of mathematics* (EMS Press, July 2023), pp. 771–823.
- [335] F. Dominguez, J. Unger, M. Traube, B. Mant, C. Ertler, and W. Lechner, “Encoding-independent optimization problem formulation for quantum computing”, *Frontiers in Quantum Science and Technology* **2**, 10.3389/frqst.2023.1229471 (2023).
- [336] Ö. Salehi, A. Glos, and J. A. Miszczak, “Unconstrained binary models of the travelling salesman problem variants for quantum optimization”, *Quantum Information Processing* **21**, 10.1007/s11128-021-03405-5 (2022).
- [337] H. F. Trotter, “On the product of semi-groups of operators”, *Proceedings of the American Mathematical Society* **10**, 545–551 (1959).
- [338] S. H. Sack and M. Serbyn, “Quantum annealing initialization of the quantum approximate optimization algorithm”, *Quantum* **5**, 491 (2021).
- [339] S. Blanes, F. Casas, and A. Murua, “Splitting methods for differential equations”, *Acta Numerica* **33**, 1–161 (2024).

- [340] T. N. Ikeda, A. Abrar, I. L. Chuang, and S. Sugiura, “Minimum Trotterization formulas for a time-dependent Hamiltonian”, *Quantum* **7**, 1168 (2023).
- [341] L. Zhou, S.-T. Wang, S. Choi, H. Pichler, and M. D. Lukin, “Quantum approximate optimization algorithm: performance, mechanism, and implementation on near-term devices”, *Physical Review X* **10**, 021067 (2020).
- [342] S. Boulebnane, J. Sud, R. Shaydulin, and M. Pistoia, *Quantum approximate optimization algorithm in finite size and large depth and equivalence to quantum annealing*, 2025.
- [343] P. Díez-Valle, F. J. Gómez-Ruiz, D. Porras, and J. J. García-Ripoll, *Universal resources for qaoa and quantum annealing*, 2025.
- [344] M. Larocca, P. Czarnik, K. Sharma, G. Muraleedharan, P. J. Coles, and M. Cerezo, “Diagnosing barren plateaus with tools from quantum optimal control”, *Quantum* **6**, 824 (2022).
- [345] A. Galda, E. Gupta, J. Falla, X. Liu, D. Lykov, Y. Alexeev, and I. Safro, “Similarity-based parameter transferability in the quantum approximate optimization algorithm”, *Frontiers in Quantum Science and Technology* **2**, 10.3389/frqst.2023.1200975 (2023).
- [346] K. Khodjasteh and D. Lidar, “Fault-tolerant quantum dynamical decoupling”, in 2005 quantum electronics and laser science conference, QELS-05 ().
- [347] C. H. Valahu, A. M. Lawrence, S. Weidt, and W. K. Hensinger, “Robust entanglement by continuous dynamical decoupling of the j-coupling interaction”, *New Journal of Physics* **23**, 113012 (2021).
- [348] W. Morong, K. Collins, A. De, E. Stavropoulos, T. You, and C. Monroe, “Engineering dynamically decoupled quantum simulations with trapped ions”, *PRX Quantum* **4**, 010334 (2023).
- [349] P. Barthel, P. H. Huber, J. Casanova, I. Arrazola, D. Niroomand, T. Sriarunothai, M. B. Plenio, and C. Wunderlich, “Robust two-qubit gates using pulsed dynamical decoupling”, *New Journal of Physics* **25**, 063023 (2023).
- [350] M. Nünnerich, D. Cohen, P. Barthel, P. H. Huber, D. Niroomand, A. Retzker, and C. Wunderlich, *Fast, robust and laser-free universal entangling gates for trapped-ion quantum computing*, 2024.
- [351] G. Quiroz and D. A. Lidar, “High-fidelity adiabatic quantum computation via dynamical decoupling”, *Physical Review A* **86**, 042333 (2012).
- [352] S. Blanes, F. Casas, J. Oteo, and J. Ros, “The magnus expansion and some of its applications”, *Physics Reports* **470**, 151–238 (2009).
- [353] P. Huber, *Multi-qubit gates in a trapped-ion quantum computer*, 2024.
- [354] D. B. Percival, “Simulating gaussian random processes with specified spectra”, *Computing Science and Statistics*, 24, pp. 534–538 (1992).
- [355] D. Sherrington and S. Kirkpatrick, “Solvable model of a spin-glass”, *Physical Review Letters* **35**, 1792–1796 (1975).
- [356] D. Panchenko, *The sherrington-kirkpatrick model* (Springer New York, 2013).
- [357] E. Boros and P. L. Hammer, “Pseudo-boolean optimization”, *Discrete Applied Mathematics* **123**, 155–225 (2002).

- [358] T. D. Hansen, H. Kaplan, O. Zamir, and U. Zwick, “Faster k -sat algorithms using biased-pps_z”, in Proceedings of the 51st annual acm sigact symposium on theory of computing, STOC ’19 (June 2019), pp. 578–589.
- [359] P. Rakyta and Z. Zimborás, “Approaching the theoretical limit in quantum gate decomposition”, 10.48550/ARXIV.2109.06770 (2021).
- [360] L. Madden and A. Simonetto, “Best approximate quantum compiling problems”, ACM Transactions on Quantum Computing **3**, 1–29 (2022).
- [361] M. M. Rams, M. Mohseni, and A. d. Campo, “Inhomogeneous quasi-adiabatic driving of quantum critical dynamics in weakly disordered spin chains”, New Journal of Physics **18**, 123034 (2016).
- [362] M. Mohseni, J. Strumpfer, and M. M. Rams, “Engineering non-equilibrium quantum phase transitions via causally gapped hamiltonians”, New Journal of Physics **20**, 105002 (2018).
- [363] Y. Susa, Y. Yamashiro, M. Yamamoto, I. Hen, D. A. Lidar, and H. Nishimori, “Quantum annealing of the p-spin model under inhomogeneous transverse field driving”, Physical Review A **98**, 042326 (2018).
- [364] M. Dadgar and C. L. Baldwin, “Anomalously slow dynamics of inhomogeneous quantum annealing”, Phys. Rev. A **111**, 062601 (2025).
- [365] M. Troyer and U.-J. Wiese, “Computational complexity and fundamental limitations to fermionic quantum monte carlo simulations”, Physical Review Letters **94**, 170201 (2005).
- [366] M. Marvian, D. A. Lidar, and I. Hen, “On the computational complexity of curing non-stoquastic hamiltonians”, Nature Communications **10**, 10.1038/s41467-019-09501-6 (2019).
- [367] J. Klassen, M. Marvian, S. Piddock, M. Ioannou, I. Hen, and B. M. Terhal, “Hardness and ease of curing the sign problem for two-local qubit hamiltonians”, SIAM Journal on Computing **49**, 1332–1362 (2020).
- [368] S. Bravyi, A. J. Bessen, and B. M. Terhal, *Merlin-Arthur games and stoquastic complexity*, arXiv:quant-ph/0611021, 2006.
- [369] S. Bravyi and B. M. Terhal, “Complexity of stoquastic frustration-free Hamiltonians”, SIAM Journal on Computing **39**, 1462–1485 (2010).
- [370] M. B. Hastings, “The power of adiabatic quantum computation with no sign problem”, Quantum **5**, 597 (2021).
- [371] A. Gilyén, M. B. Hastings, and U. Vazirani, “(sub)exponential advantage of adiabatic quantum computation with no sign problem”, in Proceedings of the 53rd annual acm sigact symposium on theory of computing (stoc) (2021), pp. 1357–1369.
- [372] B. Seoane and H. Nishimori, “Many-body transverse interactions in the quantum annealing of the p-spin ferromagnet”, Journal of Physics A: Mathematical and Theoretical **45**, 435301 (2012).
- [373] Y. Seki and H. Nishimori, “Quantum annealing with antiferromagnetic transverse interactions for the Hopfield model”, Journal of Physics A: Mathematical and Theoretical **48**, 335301 (2015).
- [374] E. J. Crosson and D. A. Lidar, “Prospects for quantum enhancement with diabatic quantum annealing”, Nature Reviews Physics **3**, 466 (2021).

-
- [375] R. Ghosh, L. A. Nutricati, N. Feinstein, P. A. Warburton, and S. Bose, “Exponential speed-up of quantum annealing via n-local catalysts”, (2024).
 - [376] N. Feinstein, I. Shalashilin, S. Bose, and P. A. Warburton, “Robustness of diabatic enhancement in quantum annealing”, *Quantum Science and Technology* **10**, 025011 (2025).
 - [377] E. Crosson, T. Albash, I. Hen, and A. P. Young, “De-signing Hamiltonians for quantum adiabatic optimization”, *Quantum* **4**, 334 (2020).
 - [378] K. Takada, S. Sota, S. Yunoki, B. Pokharel, H. Nishimori, and D. A. Lidar, “Phase transitions in the frustrated ising ladder with stoquastic and nonstoquastic catalysts”, *Phys. Rev. Res.* **3**, 043013 (2021).
 - [379] P. Hauke, L. Bonnes, M. Heyl, and W. Lechner, “Probing entanglement in adiabatic quantum optimization with trapped ions”, *Frontiers in Physics* **3**, 10.3389/fphy.2015.00021 (2015).
 - [380] G. C. Santra, S. S. Roy, D. J. Egger, and P. Hauke, “Genuine multipartite entanglement in quantum optimization”, *Physical Review A* **111**, 022434 (2025).
 - [381] C. Capecchi, G. C. Santra, A. Bottarelli, E. Tirrito, and P. Hauke, *Role of nonstabilizerness in quantum optimization*, 2025.
 - [382] S. Bravyi and A. Kitaev, “Universal quantum computation with ideal clifford gates and noisy ancillas”, *Phys. Rev. A* **71**, 022316 (2005).
 - [383] J. K. Stockton, J. M. Geremia, A. C. Doherty, and H. Mabuchi, “Characterizing the entanglement of symmetric many-particle spin- $\frac{1}{2}$ systems”, *Phys. Rev. A* **67**, 022112 (2003).
 - [384] T. Haug and L. Piroli, “Stabilizer entropies and nonstabilizerness monotones”, *Quantum* **7**, 1092 (2023).
 - [385] G. Passarelli, R. Fazio, and P. Lucignano, “Nonstabilizerness of permutationally invariant systems”, *Phys. Rev. A* **110**, 022436 (2024).
 - [386] P. S. Tarabunga, E. Tirrito, T. Chanda, and M. Dalmonte, “Many-body magic via Pauli-Markov chains—from criticality to gauge theories”, *PRX Quantum* **4**, 040317 (2023).