

Robust Coherence and Entanglement Creation in Trapped Ions

Jake Lishman

April 2022

*Thesis submitted in partial fulfilment of the requirements of the degree of
Doctor of Philosophy in the Department of Physics of Imperial College London.*

Copyright

© 2022 Jake Lishman. The copyright of this thesis rests with the author. Unless otherwise indicated, its contents are licensed under a [Creative Commons Attribution 4.0 International Licence](#) (CC BY). Under this licence, you may copy and redistribute the material in any medium or format for both commercial and non-commercial purposes. You may also create and distribute modified versions of the work. This on the condition that you credit the author.

When reusing or sharing this work, ensure you make the licence terms clear to others by naming the licence and linking to the licence text. Where a work has been adapted, you should indicate that the work has been changed and describe those changes. Please seek permission from the copyright holder for uses of this work that are not included in this licence or permitted under UK Copyright Law.

This thesis is my own work. Contributions by other people are attributed and referenced. Many of the results in this thesis can be found in the published works:

- Jake Lishman and Florian Mintert, *Trapped-Ion Entangling Gates Robust Against Qubit Frequency Errors*, [Physical Review Research](#) **2**, 033117 (2020);
- Mahdi Sameti, Jake Lishman, and Florian Mintert, *Strong-Coupling Quantum Logic of Trapped Ions*, [Physical Review A](#) **103**, 052603 (2021); and
- Ollie Corfield*, Jake Lishman*, Chungsun Lee, Jacopo Mosca Toba, George Porter, Johannes Heinrich, Simon Webster, Florian Mintert, and Richard Thompson, *Certifying Multilevel Coherence in the Motional State of a Trapped Ion*, [PRX Quantum](#) **2**, 040359 (2021).

Several figures in this document are derived from these publications, used with permission from the American Physical Society. Figures 4.1, 4.2, 5.1 to 5.3 and 6.2 to 6.4 have very minor modifications—to match the style of this document—from their counterparts in Corfield *et al.*²⁹ (those in chapter 4), Lishman and Mintert⁷² (chapter 5) and Sameti, Lishman, and Mintert¹⁰⁴ (chapter 6). These were originally authored by me, with input from my co-authors, and are largely regenerated from the same code. Figures 4.3 to 4.5 are my own versions of figures from Corfield *et al.*²⁹, and fig. 6.1 is based on a figure from Sameti, Lishman, and Mintert¹⁰⁴.

*Shared the lead author credit.

Abstract

Both entanglement and coherence are key resources for all applications of quantum technologies, from the well-known efforts to create a quantum computer, to research into thermodynamic work extraction. Trapped ions are one of the leading platforms for scalable quantum computing, as the site of many of the earliest quantum logic gates, and now boast the current highest-fidelity gates and longest coherence times of their qubits. This thesis presents three strands of work surrounding the creation, manipulation and verification of coherence and entanglement in trapped ions.

Coherence is classified into differing ranks, to better represent the structure of multiple-component superpositions. A certifier for these different levels, analogous to an entanglement witness, is derived from a one-dimensional interference pattern in a generalisation of the Ramsey scheme. This metric cannot produce false positives for high-order coherence, even when the coherence basis cannot be measured directly. It requires significantly fewer experimental resources than alternate schemes that have been proposed, and a demonstration in the motional mode of a single trapped ion is presented, verifying that 3-coherence was created.

The Mølmer–Sørensen Bell-state-creation gate in trapped ions is then examined, and its principal sources of frequency errors investigated. A multi-tone extension of the gate is presented, which is numerically optimised to make its entanglement generation robust against errors in the qubit and driving frequencies. This analysis produces a gate that is specifically optimised for the estimated error distributions of the target experiment.

Finally, the same Mølmer–Sørensen gate is taken outside the weak-coupling approximation in which it has hitherto been confined. A new method of perturbative expansion is introduced and used to calculate functional constraints on the applied driving fields that can be satisfied to cancel unwanted non-linear terms from the dynamics order-by-order. This new strategy removes a previously fundamental limitation on the speed of trapped-ion entangling gates, and severely relaxes the cooling requirements on the motional modes.

Acknowledgements

I leave Imperial with nothing but positive experiences to look back on, and I'm grateful to all those with whom I worked in one capacity or another. There are some people deserving of my particular thanks.

My supervisor FLORIAN MINTERT was unfailing in his support for me even after I crashed into him while cycling to Kingston, and I never stopped being amazed at his ability to discuss whatever physics I had on my mind when I knocked on his door, without a second's preparation. My co-supervisor RICHARD THOMPSON's caring approach both scientifically and emotionally always gave the encouragement I needed to succeed, no matter what predicament I got myself into.

I would not have stayed sane without FRED SAUVAGE to work alongside and thrash me at squash, although I did have to wage constant war against the hordes of papers invading my desk. OLLIE CORFIELD welcomed me into the ion-trap lab, and my fondest memory was our first clumsy attempt at the experiment that three years later became chapter 4, once we'd figured out what we were doing and evicted me back to the theory office. SIMON WEBSTER was my adoptive postdoc, always willing to answer my naïve theorist's questions about the lab, and who motivated me to keep running both literally and figuratively.

My partner MARTA SEGIT showed me how to celebrate my achievements, and helped me when I was panicking about having to give my first ever technical presentation. I repaid her by becoming nocturnal in my final rush to write this thesis and seeing her only in passing for a month, despite living in the same house. She has been the greatest model of support in the six-and-a-half years I have known her, and I could not have completed my degree without her.

Finally, my father GRAHAM LISHMAN. If Marta got me over the finish line, it was my Dad who made it possible for me to even begin. I write this now two decades on from the night Dad taught me to solve simple simultaneous equations after I got stuck on a problem at school. While the maths lessons stopped not long after, the ways he consistently praised effort and perseverance instead of simple results and showed me that I was loved every day of my life are why I am here today. I can only hope to be as good a parent as you have been.

—JAKE

Contents

I	Background	9
1	Introduction	10
1.1	Outline	11
2	Quantum Information	14
2.1	Basic definitions	14
2.2	Qubits and harmonic-oscillator systems	16
2.3	Measurements	18
2.4	Time evolution	19
3	Ion-Trap Quantum Computing	23
3.1	Qubit encodings	24
3.2	Trapped-ion dynamics	26
3.3	Ion–laser interactions	30
3.3.1	General Hamiltonian	30
3.3.2	Sideband transitions	33
3.4	Mølmer–Sørensen gate	36
II	New Research	40
4	Certification of Higher-Order Coherence	41
4.1	Quantum coherence	42
4.2	Interference-pattern methods	44
4.3	General measurements	48
4.3.1	Analytic threshold for 3-coherence	48
4.3.2	Numeric evaluation of thresholds	50
4.4	State-creation sequences	52
4.5	Measurement-mapping sequences	55
4.6	Statistics of the certifier	60

Contents

4.7	Experimental realisation	64
4.7.1	State creation	65
4.7.2	Coherence certification	68
4.8	Conclusion	72
5	Robust Entangling Gates	74
5.1	Model	75
5.2	Optimisation	78
5.2.1	Parametrisation	79
5.2.2	Reduction of dimensionality	80
5.2.3	Quadrature	82
5.3	Results	84
5.4	Outlook	89
6	Entangling Gates With Strong Coupling	91
6.1	Non-linear ion–motion interactions	92
6.2	Calculating solutions	98
6.2.1	Finding constraints	98
6.2.2	Solving constraints	99
6.2.3	Example solutions	101
6.3	Simulation results	103
6.4	Additional motional modes	106
6.5	Outlook	108
7	Conclusion	110

List of Figures

3.1	Trapped-ion qubit encodings and addressing schemes	24
3.2	Energy-level structure of $^{40}\text{Ca}^+$	26
3.3	First-order sideband transitions in a single trapped ion	34
3.4	Mølmer-Sørensen gate energy-level scheme	37
4.1	Creation of an arbitrary motional superposition	53
4.2	PDFs of biased and corrected estimators	63
4.3	Dynamics of the motional blue sideband transition	66
4.4	Interference patterns for two- and three-element superpositions	70
4.5	Interference patterns of a four-element motional superposition	70
5.1	Energy-level diagram of two co-trapped ions	76
5.2	Fidelity properties of the robust multi-tone Mølmer-Sørensen gate	86
5.3	Phase-space paths during the robust multi-tone Mølmer-Sørensen gate	88
6.1	Energy-level diagram of two co-trapped ions with motion	93
6.2	Gate infidelity of the strongly coupled Mølmer-Sørensen gate	103
6.3	Performance of the strongly coupled hot Mølmer-Sørensen gate	104
6.4	Phase-space paths of the strongly coupled Mølmer-Sørensen gate	105

List of Tables

3.1	Equilibrium positions of ions in a linear chain	28
3.2	Joint normal modes of motion of ions in a chain	29
4.1	Creation and mapping sequences for $(g, 0\rangle + g, 1\rangle)/\sqrt{2}$	59
4.2	Creation and mapping sequences for $(g, 0\rangle + g, 1\rangle + g, 2\rangle)/\sqrt{3}$	59
4.3	Creation and mapping sequences for $(g, 0\rangle + g, 1\rangle + g, 2\rangle + g, 3\rangle)/2$...	59
4.4	Fit parameters after superposition creation	67
4.5	Fitted motional populations after superposition creation	67
5.1	Error equivalence rules by frame transformations	81
5.2	Weight functions and associated orthogonal polynomials	83
5.3	Optimised driving schemes for multi-tone Mølmer-Sørensen gates	85

Part I

Background

Chapter 1

Introduction

The last decade has seen the beginnings of large-scale commercialisation of quantum computing, although practical advantages over classical computers still hover out of reach. The field arose forty years ago out of discussions about quantum simulations of nature³⁹, and possible quantum descriptions of Turing machines^{9,10}. Within a decade, this developed into prospective models of a *universal quantum computer*³² and the first problem for which quantum systems offered an improvement in computational complexity over the best classical solution³³. Non-academic interest in quantum information processing intensified after the publication of Shor’s algorithm in 1994¹¹², which demonstrated that the integer-factorisation problem, a bedrock of modern asymmetric cryptography, could be solved in a time polynomial in the number of bits rather than the sub-exponential asymptotic scaling of the most efficient known classical approach.

For time being, our bank transactions and text messages remain safe. Since the turn of the millennium, several groups have successfully determined that $15 = 3 \times 5$ using various quantum systems^{67,74,126}, and in the last year, one group found that $21 = 3 \times 7$ on a commercial superconducting-qubit IBM machine¹¹³. Of course, this is a somewhat facetious point; the immediate goals are to show that the methods are viable in the current *noisy intermediate-scale* regimes of quantum devices. Still, though, five-bit integers are a far cry from the 4096-bit products that have become the standard for RSA public-key systems. Current implementations are simply unable to fabricate or control millions of qubits at the operational tolerances required for these applications. Much of this is due to errors in the logical operations. While modern CPUs are—cosmic rays aside—essentially error-free in actual usage, this is not true of their quantum counterparts.

The achievable interaction fidelities have steadily improved since the first two-qubit logic gate in trapped ions⁸³. The current state of the art remains in this same setting, with two-qubit gates now recorded at close to 99.99% fidelity^{4,40}. This is not the only medium for quantum computing, however. Early quantum algorithms were demonstrated in nuclear magnetic resonance systems^{24,62}, and

in the intervening years, logic gates have been demonstrated in superconducting qubits¹³⁵, linear photonic systems⁶⁵, neutral atoms¹⁵, and other systems.

The requirements for a scalable quantum information processor are generally agreed upon³⁵. The five operations are: a set of well defined, addressable qubits; a reliable method to prepare known quantum states; decoherence times much longer than gate operations; access to a universal set of quantum logic gates; and a method of reading out the state of a qubit in some basis. While modern gate fidelities are high, decoherence processes pose limitations on how much further they can realistically be pushed on a large scale. With known quantum algorithms requiring many thousands or millions of gates, even the current state-of-the-art infidelities of around one part in ten thousand are insufficient. Instead, theoretical work into quantum error correction has shown that *fault-tolerant* quantum computation can still be achieved with fidelities on the order of 99–99.9%^{36,100}. This is effectively reducing requirements on fidelities by increasing the number of qubits needed, elevating further the importance of scalability. The new infidelity goals are within reach for small numbers of qubits, but there are massive hurdles to overcome in experimental control and qubit isolation as register sizes increase.

Of the candidates, trapped-ion and superconducting qubits currently seem the most likely to successfully scale in the near term. Both of these have serious commercial backing: trapped ions by Honeywell⁹³ and IonQ¹⁴, and superconducting qubits by Google³ and IBM¹⁴¹. Trapped ions have the better fidelities and state lifetimes relative to their gate speeds, but the absolute gate speeds of superconducting qubits are orders of magnitude faster and their fabrication can build on the back of existing silicon technology. Both remain highly susceptible to many decoherence processes from the environment. This work focusses entirely on trapped ions, aiming to move the technology closer to satisfying all of the quantum-computing requirements completely.

1.1 Outline

The second part of this work contains three main areas of novel research, all linked by the goal of enabling more robust coherence and entanglement generation in noisy trapped-ion systems, key ingredients of creating a large-scale quantum information processor. These are not presented in chronological order, but instead progress from dealing with single ions, then to two-ion entanglement in a weakly coupled regime, and finish on a method for implementing two-qubit gates in non-linear regimes of strongly coupled interactions with multiple hot motional modes. All

three of these projects led to publications^{29,72,104}, with chapters 4 and 6 both being collaborative efforts with other researchers. Contributions from my co-authors are clearly stated at the beginning of each chapter, and I was heavily involved with all of the work that I have described in the rest of this thesis.

As with any PhD thesis in modern physics, this work stands on the shoulders of programmatic giants. The subsequent text will not delve into the minutiae of any code, but I would be remiss were I not to mention those without whom I could never have produced this thesis. All the results of chapters 4 to 6 used significant computational resources to achieve, provided by the Research Computing Service at Imperial College London⁵⁰. The majority of the programming here was built on the NumPy⁴⁷ and SciPy¹²⁸ packages. The library QuTiP⁶¹ was used for integration of arbitrary time-dependent systems, a project of which I became a principal contributor and then maintainer over the course of my degree.

Chapter 4 describes a generalisation of Ramsey-like interference-pattern experiments to robustly certify the presence of multilevel coherence in the motional state of a single trapped ion. This can never return false positives, despite the motional mode being inaccessible to measurement. In conjunction with the experimental group at Imperial, we implemented this scheme in a real-world system, and unambiguously verified that we had created three-level coherence.

My first research project is presented in chapter 5, where we perform numerical optimisations of a multitone extension to the Mølmer–Sørensen entangling gate, to make it resilient against fluctuations and miscalibrations of the individual qubit frequencies. Our simulations indicate potential order-of-magnitude improvements in the infidelity of the gate at the error-correction threshold, and the methods used can efficiently handle any calibrated error model.

Finally, chapter 6 describes a systematic method for moving trapped-ion gates outside of the weakly coupled, linear regime they have hitherto been confined within. This involves driving higher-order motional transitions with very simple control fields that any current ion-trap group could easily implement with their existing hardware. We illustrate a perturbative expansion to determine the gate dynamics order-by-order of the coupling strength for a general non-linear interaction, and derive a series of functional constraints on the driving profiles that, when satisfied, allow the gate to be decoupled from the motion to ever higher orders. We show two explicit solutions to these, improving the infidelity scaling of the gate by several orders. This reduces the previously fundamental limitations on gate fidelity by a factor of 2000 with the simplest extension, and allows gates to operate without expensive sideband cooling cycles. For an outlook, we sketch out potential

procedures for cancelling non-linearities from spectating motional modes and the steps needed to incorporate existing dynamical-decoupling or robust-gate schemes into this new framework.

The rest of the first part of this work is devoted to the other introductory material needed for the new research. This includes the basic definitions and mathematical techniques of quantum mechanics and information in chapter 2, and then the relevant physics of trapped ions in chapter 3.

Chapter 2

Quantum Information

The mathematical language of quantum mechanics is physicists' dialect of linear algebra. We are exclusively concerned with complex Hilbert spaces: vector spaces equipped with an inner product. This chapter is primarily intended as a reference for the notation and terminology that will be used in the rest of the thesis. A far more complete introduction to quantum mechanics, quantum information theory, and their mathematical backing may be found in the venerable Nielsen and Chuang⁸⁸.

2.1 Basic definitions

Pure quantum states are elements of a vector space \mathcal{V} over the field of complex numbers \mathbb{C} . We write a state in *bra-ket* notation as a *ket* $|\psi\rangle$, where ψ is an identifier rather than necessarily having any mathematical properties. As would be expected from the properties of a vector space, such systems can be in linear superpositions of states $a|\psi\rangle + b|\phi\rangle$, where a and b are complex numbers. These superpositions are the foundation of quantum coherence, which is responsible for many of the counterintuitive predictions of quantum mechanics in which a state interferes with itself⁵³.

An *operator* \hat{A} within the same Hilbert space is a mapping $\mathcal{V} \rightarrow \mathcal{V}$, which takes one state to another by acting on it as $\hat{A}|\psi\rangle = |\phi\rangle$. This thesis will not stray beyond *linear* operators, so $\hat{A}(a|\psi\rangle + b|\phi\rangle) = a\hat{A}|\psi\rangle + b\hat{A}|\phi\rangle$ for all complex scalars a and b , and all states $|\psi\rangle$ and $|\phi\rangle$. Operators do not, in general, commute—that is $\hat{A}\hat{B} \neq \hat{B}\hat{A}$ for most \hat{A} and \hat{B} . The *commutator*, defining the difference between the two, is written $[\hat{A}, \hat{B}] = \hat{A}\hat{B} - \hat{B}\hat{A}$. Functions of operators are defined by means of their power series. For example, the exponential of an operator can be written as

$$\exp(\hat{A}) = 1 + \hat{A} + \frac{1}{2!}\hat{A}^2 + \frac{1}{3!}\hat{A}^3 + \dots, \quad (2.1)$$

where loose scalars are implicitly multiplied by the suitable identity operator.

The inner product between two states is written as $\langle\phi|\psi\rangle$, where in contrast

to mathematical notation, the operation is linear in $|\psi\rangle$ and conjugate-linear in $\langle\phi|$. By the time we reach the new results in this thesis, we will be dealing solely with normalised states, such that $\langle\psi|\psi\rangle = 1$. The inner product of a state with the output of an operator acting on a state is written $\langle\phi|\hat{A}|\psi\rangle$, and is colloquially called a *matrix element* by analogy to the matrix representation of linear algebra. In practice we rarely write states or operators in column-vector or matrix form, and we will often deal with infinite-dimensioned Hilbert spaces where this would be impractical at best.

The object $\langle\psi|$ is named a *bra* and is an element of the dual of the vector space. Functionally, it is a linear mapping $\mathcal{V} \rightarrow \mathbb{C}$ taking elements of the vector space to the complex field, and so can apply to states and cannot commute through them. The object $|\psi\rangle\langle\phi|$ is therefore an operator in its own right. Operators can also act on bras, with the definition $(\langle\phi|\hat{A})|\psi\rangle = \langle\phi|(\hat{A}|\psi\rangle) = \langle\phi|\hat{A}|\psi\rangle$ for all $|\psi\rangle$ and $|\phi\rangle$.

The *adjoint* of an operator \hat{A}^\dagger is the operator such that the inner product of $\hat{A}|\psi\rangle$ on $|\phi\rangle$ is equal to inner product of $|\psi\rangle$ on $\hat{A}^\dagger|\phi\rangle$. For convenience, although it is not strictly mathematically accurate to do so, in bra-ket notation we define $\langle\psi| = (|\psi\rangle)^\dagger$. This is a deliberate choice to simplify the inner-product notation; it is irrelevant to the physics whether $\langle\phi|\hat{A}|\psi\rangle$ came from $(|\phi\rangle)^\dagger(\hat{A}|\psi\rangle)$ or $(\hat{A}^\dagger|\phi\rangle)^\dagger|\psi\rangle$. *Hermitian* (self-adjoint) operators $\hat{A} = \hat{A}^\dagger$ are especially important in quantum mechanics, because all physical observables must be of this form. Hermitian operators have real eigenvalues, and their eigenstates form complete orthonormal bases of the respective Hilbert space. Throughout this thesis, the notation $\hat{A} + \text{H.c.}$ means $\hat{A} + \hat{A}^\dagger$, with H.c. standing for *Hermitian conjugate*.

A *unitary* operator is an operator \hat{U} such that $\hat{U}\hat{U}^\dagger = \hat{U}^\dagger\hat{U} = 1$, i.e. an operator whose adjoint is its inverse. These operators are inner-product-preserving, and consequently all valid evolutions taking one physical pure state to another are unitary, including quantum logic gates. All unitary operators can be written as $\hat{U} = \exp(i\hat{H})$ for some Hermitian operator \hat{H} . Writing \hat{H} in terms of its eigenvalue decomposition $\hat{H} = \sum_j \lambda_j |\lambda_j\rangle\langle\lambda_j|$ for orthonormal $|\lambda_i\rangle$, it is clear that that eigenvalues of this \hat{U} are $e^{i\lambda_j}$, which have unity magnitude. This form of operator frequently appears when describing evolution of a quantum system.

It is often the case that a quantum-mechanical system is subject to some classical interaction that reduces the quantum coherence and introduces in its place some classical probability of being in a particular state. These states are represented by a *density operator* $\hat{\rho}$ as

$$\hat{\rho} = \sum_j p_j |\psi_j\rangle\langle\psi_j|, \quad (2.2)$$

where the $\{p_j\}$ are classical probabilities that sum to one. The density-operator representation of a pure state has exactly one nonzero probability, and the corresponding $|\psi_j\rangle$ is the vector representation of the pure state. If an operator were applied to the state such that $|\psi_j\rangle \rightarrow \hat{A}|\psi_j\rangle$, the new density operator would be $\hat{A}|\psi_j\rangle\langle\psi_j|\hat{A}^\dagger$, illustrating that density operators evolve as $\hat{\rho} \rightarrow \hat{A}\hat{\rho}\hat{A}^\dagger$.

A quantum system can also comprise more than one separate physical system. Formally, the joint space is the tensor product of the component vector spaces $\mathcal{V}_1 \otimes \mathcal{V}_2$, which is itself a vector space. We will typically write kets in a joint space by concatenating the labels, such as $|g\rangle \otimes |g\rangle = |gg\rangle$, or by juxtaposition, such as $(|g\rangle + |e\rangle)(|g\rangle + |e\rangle) = |gg\rangle + |ge\rangle + |eg\rangle + |ee\rangle$. Operators on a joint space will similarly be juxtaposed, and if an operator is missing for a particular subspace, it is implicitly the identity. In cases where there could be ambiguity between joint operators and composed operator action, we will use the tensor-product symbol \otimes and leave composition as-is.

All states in a tensor-product space can be written as $|\psi\rangle = \sum_{j,k} c_{j,k} |j\rangle \otimes |k\rangle$, where the sum is over arbitrary orthonormal bases of the two subspaces. This decomposition is not unique; it is dependent on the bases chosen. A state is said to be *entangled* when there must be more than one nonzero element in the sum, regardless of its magnitude, while unentangled states are called *separable*. This can be extended to a hierarchical structure, accounting for the number of entangled subsystems¹²³, and there is a great body of literature considering the verification, distillation and use of entanglement as a resource^{22,37,56}.

A very similar approach allows us to define multilevel *coherence*, which is defined with respect to a particular choice of basis, rather than the fairly natural separation of physical systems used in entanglement. Given a basis $\{|j\rangle\}$, a pure state is called k -coherent if at least k basis vectors have nonzero overlap with it. For a mixed state $\hat{\rho}$ to be k -coherent, all its possible decompositions must contain at least one pure state that is k -coherent or greater. Precisely determining the level of coherence of a large system is nontrivial, even if the density operator is known exactly. This is the focus of chapter 4.

2.2 Qubits and harmonic-oscillator systems

Only two types of quantum system will be used in this thesis: qubits and quantum harmonic oscillators. We will briefly cover the notation used to work with these.

The name *qubit* is used to describe a physical object whose states reside in a two-dimensional complex Hilbert space, for example a two-level ion. We will label

the two basis states $|g\rangle$ and $|e\rangle$ (for *ground* and *excited*), to reduce confusion with the harmonic oscillator number states. In this form, the Pauli operators are

$$\hat{\sigma}_x = |e\rangle\langle g| + |g\rangle\langle e|, \quad \hat{\sigma}_y = -i|e\rangle\langle g| + i|g\rangle\langle e|, \quad \text{and} \quad \hat{\sigma}_z = |e\rangle\langle e| - |g\rangle\langle g|, \quad (2.3)$$

where $\hat{\sigma}_z$ is related to the free evolution of the system, while $\hat{\sigma}_x$ and $\hat{\sigma}_y$ are related to transitions between the two states. The Pauli operators are both Hermitian and unitary, so each of their squares is simply the identity operator. The multiplication of distinct Pauli operators satisfies $\hat{\sigma}_a \hat{\sigma}_b = i\varepsilon_{abc} \hat{\sigma}_c$ for a, b and c in $\{x, y, z\}$, where ε is the Levi-Civita symbol with parity defined by $\varepsilon_{xyz} = 1$. When dealing with coupled systems, we will also use two related operators

$$\hat{\sigma}_+ = |e\rangle\langle g| = \frac{1}{2}(\hat{\sigma}_x + i\hat{\sigma}_y) \quad \text{and} \quad \hat{\sigma}_- = |g\rangle\langle e| = \frac{1}{2}(\hat{\sigma}_x - i\hat{\sigma}_y). \quad (2.4)$$

Taken individually these are non-Hermitian and non-unitary, representing the separate excitation and de-excitation processes, and will be useful in rotating-wave approximations.

The eigensystem of a quantised harmonic oscillator is spanned by the Fock or *number* basis. We label the states with an integer n as $|n\rangle$, where $|0\rangle$ is the ground state, and so on. The principal operators when working with these states are the annihilation \hat{a} and creation \hat{a}^\dagger operators, also called *ladder* operators, that respectively remove and add a phonon of motion to the system by the relations

$$\hat{a}|n\rangle = \sqrt{n}|n-1\rangle \quad \text{and} \quad \hat{a}^\dagger|n\rangle = \sqrt{n+1}|n+1\rangle. \quad (2.5)$$

Note that $\hat{a}|0\rangle = 0$; the state cannot be lowered beyond the ground state. The two operators do not commute, but satisfy $[\hat{a}, \hat{a}^\dagger] = 1$. The Hermitian operator $\hat{a}^\dagger \hat{a}$ is named the *number* operator, as it clearly follows from eq. (2.5) that $\hat{a}^\dagger \hat{a}|n\rangle = n|n\rangle$.

Physically, the two ladder operators arise from the diagonalisation of the quantum harmonic oscillator Hamiltonian, and are defined by $\hat{a} = \hat{x} - i\hat{p}$ for nondimensionalised position \hat{x} and momentum \hat{p} . Following on, the *displacement* operator

$$\hat{D}(\alpha) = \exp(\alpha \hat{a}^\dagger - \alpha^* \hat{a}) \quad (2.6)$$

displaces a state by an amount α in phase space, where the real and imaginary components correspond to the positional and motional displacements respectively. This is unitary, but not Hermitian—instead, $\hat{D}^\dagger(\alpha) = \hat{D}(-\alpha)$, which is geometrically intuitive. We will frequently use phase space to make qualitative interpretations

of ion-trap gate operations, since it encodes the first orders of the behaviour of motional modes.

2.3 Measurements

Manipulation of quantum systems is all very well, but we cannot gain any information until we perform a measurement. These observations are non-unitary, and in general collapse a state down to subspaces associated with each possible outcome, destroying coherence. The majority of possible measurements of quantum systems are *projective*, where the possible outcomes for a particular measurement are each defined by a positive-semidefinite Hermitian operator $\{\hat{P}_j\}$ such that $\hat{P}_j\hat{P}_k = \delta_{jk}\hat{P}_j$, and $\sum_j \hat{P}_j = 1$. These requirements imply that the \hat{P}_j partition the total Hilbert space into a collection of orthogonal subspaces, and each projector can be written $\hat{P}_j = \sum_k |\psi_k^{(j)}\rangle\langle\psi_k^{(j)}|$, where the states with equal j are an orthonormal basis of the relevant subspace. The rank of such an orthogonal projector is the number of states required in its sum representation.

Real quantum systems can often perform only one type of measurement: projection onto some logical basis. There is typically one natural basis that encodes a count or choice between physical items, so the associated measurement is simply observing which level a system is in, or how many photons or phonons exist. For qubits, this logical basis is typically chosen to coincide with the Pauli Z basis, and the two measurement outcomes are associated with the operators $|g\rangle\langle g|$ and $|e\rangle\langle e|$. When investigating more information-theoretic results, however, it is appropriate to consider the more general form of measurement.

The operators associated with a measurement's outcomes being Hermitian positive-semidefinite and summing to the identity operator is analogous to a classical probability distribution. This is axiomatic. The requirement that the separate measurement outcomes are orthogonal, however, is not. Relaxing this takes us to the most general case, that of a *positive operator-value measure* (POVM). Such measures comprise a set of operators $\{\hat{A}_j\}$ that are positive semi-definite Hermitian and sum to the identity. Notably, this allows a degree of overlap between the measurement operators. Unlike for projective measurements, performing a measurement in a POVM and obtaining the outcome associated with \hat{A}_1 does not preclude a measurement on the resulting state returning \hat{A}_2 . This can be useful in situations where two states are not perfectly distinguishable, but one wishes to know with certainty which has been received at the cost of sometimes obtaining an inconclusive result. Taking a POVM comprising of scaled projectors onto the two states

orthogonal to the targets and the remainder operator needed for completeness, a result of either of the first two operators unambiguously determines the input state, while the latter gives no information. We will use this formalism later when dealing with a similar problem: unambiguously validating the presence of coherence in a system, without risk of false negatives.

2.4 Time evolution

All quantum mechanical systems obey the Schrödinger equation

$$i\hbar\partial_t|\psi\rangle = \hat{\mathcal{H}}|\psi\rangle, \quad (2.7)$$

where ∂_t is the partial derivative with respect to time and $\hbar \approx 1.05 \times 10^{-34}$ J s is the reduced Planck constant. The Hermitian operator $\hat{\mathcal{H}}$ is the Hamiltonian of the system, which corresponds to its total energy and determines its dynamics. Equivalently, we can define a unitary operator $\hat{\mathcal{U}}$ that represents the time evolution by the solution of $i\hbar\partial_t\hat{\mathcal{U}} = \hat{\mathcal{H}}\hat{\mathcal{U}}$, which allows us to write $|\psi(t)\rangle = \hat{\mathcal{U}}(t)|\psi(0)\rangle$. As with any other operator, mixed states evolve as $\hat{\mathcal{U}}\hat{\rho}\hat{\mathcal{U}}^\dagger$ in a closed system.

If the Hamiltonian commutes with itself at different times, *i.e.* $[\hat{\mathcal{H}}(t_1), \hat{\mathcal{H}}(t_2)] = 0$ for all times t_1 and t_2 , the time-evolution operator is explicitly given by

$$\hat{\mathcal{U}}(t) = \exp\left(-\frac{i}{\hbar} \int_0^t dt' \hat{\mathcal{H}}(t')\right). \quad (2.8)$$

For general time-dependent Hamiltonians, however, the time-evolution operator must be found by a general solution of the Schrödinger equation considering each of the basis states, or some perturbative expansion.

In some cases, it is possible to reduce the Hamiltonian to a solvable form by means of a frame transformation. Unitary operators can be interpreted as a mapping from one basis of a vector space to another, so taking the state $|\psi\rangle$ to $\hat{\mathcal{U}}|\psi\rangle$ is an analogue to changing the reference frame in classical mechanics. The Hamiltonian is modified by this transformation. The Schrödinger equation for these new states must be satisfied by a new Hamiltonian $\hat{\mathcal{H}}'$, so

$$\hat{\mathcal{H}}'\hat{\mathcal{U}}|\psi\rangle = i\hbar\partial_t(\hat{\mathcal{U}}|\psi\rangle) = i\hbar(\partial_t\hat{\mathcal{U}})|\psi\rangle + i\hbar\hat{\mathcal{U}}\partial_t|\psi\rangle. \quad (2.9)$$

The last term contains the time derivative from the original Schrödinger equation,

and so by inserting identity operations explicitly as $\hat{\mathcal{U}}^\dagger \hat{\mathcal{U}}$, we reach

$$\hat{\mathcal{H}}' = \hat{\mathcal{U}} \hat{\mathcal{H}} \hat{\mathcal{U}}^\dagger + i\hbar(\partial_t \hat{\mathcal{U}}) \hat{\mathcal{U}}^\dagger. \quad (2.10)$$

This is particularly useful for the common pattern of Hamiltonians that can be split into *system* and time-dependent *interaction* components as $\hat{\mathcal{H}} = \hat{\mathcal{H}}_{\text{sys}} + \hat{\mathcal{H}}_{\text{int}}$, where the system alone can be solved exactly. Typically only the effect of the interaction is interesting, and the system component can be *rotated out* by a frame transformation of the adjoint of its unitary dynamics $\hat{\mathcal{U}}_{\text{sys}}^\dagger$. For a time-invariant system Hamiltonian, this is explicitly

$$\hat{\mathcal{H}}' = e^{i\hat{\mathcal{H}}_{\text{sys}}t/\hbar} \hat{\mathcal{H}}_{\text{int}} e^{-i\hat{\mathcal{H}}_{\text{sys}}t/\hbar}. \quad (2.11)$$

A few exponentials can be evaluated exactly via their power series into simple sums of operators, such as exponentials of Pauli operators. More commonly, the exponentials must be massaged into a more convenient form. If \hat{A} and \hat{B} do not commute, then $e^{\hat{A}+\hat{B}} \neq e^{\hat{A}}e^{\hat{B}}$. Instead, one relates $e^{\hat{A}}e^{\hat{B}}$ to a single exponential $e^{\hat{C}}$ by the Baker–Campbell–Hausdorff formula¹⁶:

$$\hat{C} = \hat{A} + \hat{B} + \frac{1}{2}[\hat{A}, \hat{B}] + \frac{1}{12}([\hat{A}, [\hat{A}, \hat{B}]] - [\hat{B}, [\hat{A}, \hat{B}]]) + \dots \quad (2.12)$$

A formal proof of this requires too much Lie theory to be worth reproducing here, but can be found in the works of Bonfiglioli and Fulci¹⁶, and Hall⁴⁶. We will make much use of a result following from this formula, that

$$e^{\hat{A}}\hat{B}e^{-\hat{A}} = \hat{B} + [\hat{A}, \hat{B}] + \frac{1}{2!}[\hat{A}, [\hat{A}, \hat{B}]] + \frac{1}{3!}[\hat{A}, [\hat{A}, [\hat{A}, \hat{B}]]] + \dots, \quad (2.13)$$

of which the utility for unitary frame transformations is obvious.

When the solution of a Hamiltonian is not analytically tractable with exact frame-transformation methods, it becomes useful to pursue perturbative expansions. The Magnus expansion considers the solution of the operator Schrödinger equation of the form $\hat{\mathcal{U}}(t) = \exp \hat{M}(t)$ for some $\hat{M} = \sum_j \hat{M}_j$, where the successive terms are associated with increasing order. By inspection of the differential equation and the solution ansatz, we can recast the problem to

$$(\partial_t e^{\hat{M}}) e^{-\hat{M}} = -\frac{i}{\hbar} \hat{\mathcal{H}}. \quad (2.14)$$

Similarly to before, we will defer to prior work for a complete description⁷⁵. Qual-

itatively, one evaluates the derivative of the exponential map to find an infinite series of the form

$$\partial_t \hat{M} = c_1 \hat{\mathcal{H}} + c_2 [\hat{\mathcal{H}}, \hat{M}] + c_3 [[\hat{\mathcal{H}}, \hat{M}], \hat{M}] + \dots \quad (2.15)$$

Direct integration yields a recursive definition of \hat{M} , from which we identify the Magnus terms:

$$\begin{aligned} \hat{M}_1 &= -\frac{i}{\hbar} \int_0^t dt_1 \hat{\mathcal{H}}(t_1) \\ \hat{M}_2 &= -\frac{1}{2\hbar^2} \int_0^t dt_1 \int_0^{t_1} dt_2 [\hat{\mathcal{H}}(t_1), \hat{\mathcal{H}}(t_2)] \\ \hat{M}_3 &= \frac{i}{6\hbar^3} \int_0^t dt_1 \int_0^{t_1} dt_2 \int_0^{t_2} dt_3 \left([\hat{\mathcal{H}}(t_1), [\hat{\mathcal{H}}(t_2), \hat{\mathcal{H}}(t_3)]] + [\hat{\mathcal{H}}(t_3), [\hat{\mathcal{H}}(t_2), \hat{\mathcal{H}}(t_1)]] \right) \\ &\vdots \end{aligned} \quad (2.16)$$

The Magnus expansion is preferable to many other possibilities as it maintains the unitarity of the operator even when the series is truncated. We will visit alternative perturbative expansions, which share this property, in chapter 6 to evaluate and control complex dynamics in two-qubit operations in trapped ions.

The Schrödinger equation is strictly accurate for all (non-relativistic) quantum systems, but relies on all subspaces being explicitly accounted for. In general, a system will undergo unwanted interactions with a much larger environment. This describes an *open* quantum system, as opposed to the *closed* systems hitherto considered. There are several formalisms for investigating these behaviours¹⁸, but our needs on this front are rather limited. We will need only consider weak perturbations, where we can assume the Born–Markov conditions that the system and environment are weakly coupled and remain separable at all times, and any correlations within the environment decay much quicker than any in the system.

Under this approximation, the time evolution of the density operator for only the system of interest can be described by the Lindblad master equation⁷⁸:

$$\partial_t \hat{\rho} = -\frac{i}{\hbar} [\hat{\mathcal{H}}, \hat{\rho}] + \frac{1}{2} \sum_j (2\hat{L}_j \hat{\rho} \hat{L}_j^\dagger - \hat{L}_j^\dagger \hat{L}_j \hat{\rho} - \hat{\rho} \hat{L}_j^\dagger \hat{L}_j). \quad (2.17)$$

The $\{\hat{L}_j\}$ are *collapse* or *dissipation* operators, which represent the effects of the environment on the system. These operators must preserve the *completely positive trace-preserving* properties of the evolution. Complete positivity is the requirement that the operation, when applied to only a subsystem of a quantum state in a larger Hilbert space, maintains the positive-semidefinite nature of the whole density

matrix. This applies a condition on the collapse operators that they are *bounded*, that is that each has a finite maximum of $\langle \psi | \hat{L}_j^\dagger \hat{L}_j | \psi \rangle$ over all normalised vectors $\{|\psi\rangle\}$ in the relevant Hilbert space.

The only effect that we shall consider in this framework is motional dephasing, associated with $\hat{a}^\dagger \hat{a}$; in the trapped-ion system in use at Imperial, the time scales of qubit noise processes are far longer than those of ion motion. This operator is not strictly bounded; in an infinite-dimensional Hilbert space, its eigenvalues are the set of natural numbers. For our purposes, we will be able to consider only a finite subspace up to some maximal motional state that never becomes populated. Within this related Hilbert space, the corresponding truncated operator is bounded.

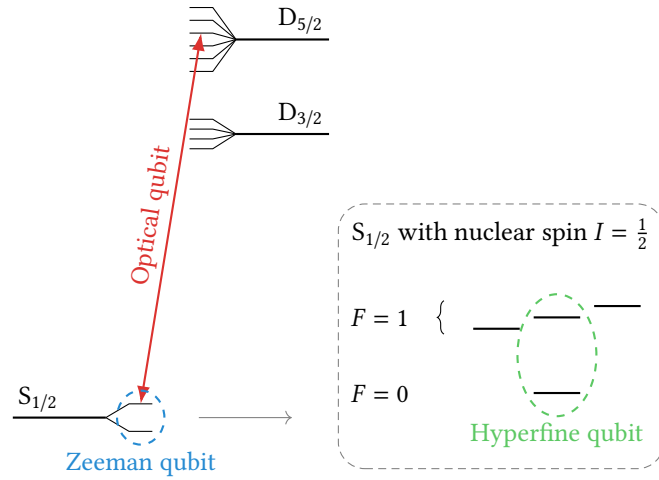
Chapter 3

Ion-Trap Quantum Computing

The electronic states of single ions are natural choices for encoding quantum information. They are clearly at the quantum scale, yet an appropriate choice can provide very long lifetimes and fabrication of each qubit is trivially reproducible; any two ions of the same charge and isotope are guaranteed to have precisely equal energy-level structures. Ions can be isolated and confined with simple electromagnetic control fields, and the Coulomb interaction between them will keep them sufficiently spatially separated that they can be individually addressed. Depending on which electronic states are used for the qubits, this may be with lasers, or with microwaves if an external field is applied to modify the transition frequencies of different qubits. The motion of the ions is coupled via the same Coulomb interaction, which once cooled into the quantum regime can be used as a communication bus between all the ions in the same trap.

Trapping is principally done with an electrostatic field to limit motion axially, and either an oscillating electric field or a static magnetic field to provide the perpendicular confinement. These two types of trap are called, respectively, linear rf or Paul traps⁹¹ and Penning traps³¹. In an idealised world with perfect experimental control, total isolation from all environmental factors, and unlimited laboratory space and budget, the two traps provide identical platforms for quantum information processing. Returning to reality, however, linear rf traps are the most common choice for quantum information processing due to simpler optical access, greater control over the radial confinement, and the non-necessity of large magnetic fields. The current highest-reported two-qubit average gate fidelities were achieved with ions in a linear rf trap, at slightly over 99.9%^{4,40}, although these results are now over five years old. This safely reaches the fidelities required for error-corrected quantum computing¹³, and more recent work out of the same groups has focussed more on performing gates at high speeds¹⁰⁷, at lower powers with microwaves¹¹⁷, and with mixed ion species⁵⁸.

FIGURE 3.1—The principal choices of encodings for qubits in trapped ions and the two major addressing schemes. The two qubit states are located in different Zeeman sublevels of the electronic ground state, different hyperfine levels of the ground-state manifold, or in an optical setup with one qubit in the ground state and the other in a metastable state in a different orbital. Frequency separations for these are on the orders of 10 MHz, 1 GHz and 100 THz respectively. Optical qubits are driven by high-powered lasers as simple two-level systems. Hyperfine and Zeeman qubits are most commonly addressed in a Raman configuration forming a lambda system, though microwave sources can theoretically drive hyperfine qubits in a single-photon mode.



3.1 Qubit encodings

In principle, a qubit can be encoded in any two distinct energy levels of a trapped ion. In practice, both of the lifetimes must be long enough to carry out all the required operations while still permitting controlled transitions, and it must be possible to reliably prepare the qubit in a known state and read out its new state later³⁵. Group-II ions are a common host for qubits as ionisation leaves a single valence electron, creating alkali-like states with simple energy-level structures. Access to an effective cooling transition is also an important part of choosing an ion, though not one integral to this thesis; for a deeper review, see ref. 38.

The preparation requirement typically ensures that at least one element of the ground-state manifold is part of the qubit, but the choice of the other state has more flexibility. Broadly, the three main routes in order of frequency separation are to encode the two states separated by Zeeman or hyperfine splitting, or use an *optical* qubit with the upper state in a different electronic orbital. These categories are broken down further by the addressing mode used to drive transitions between the qubit states: single-photon methods, or creating a lambda system with a third level and using a Raman configuration. Figure 3.1 illustrates these main schemes.


The easiest conceptually is to use the largest splitting as a simple two-level system addressed with a single laser, typically in the red and infrared range. These lasers were historically more available than those at the blue end of the spectrum used by other qubit schemes, and optical components remain easier to manufacture with lower relative transmission errors for longer wavelengths. Coherence times of these qubits are fundamentally limited by the lifetime of the excited state used. This all but requires the two levels to be separated by a dipole-forbidden transition, and

the necessity of driving a state change with the available laser power in a reasonable amount of time ensures that quadrupole transitions are by far the most feasible candidates. Lifetimes on the order of one second are typical for these quadrupole transitions⁵. However, to achieve coherent operations for this duration, the laser itself must have a linewidth on the order of 1 Hz; the dephasing rate is tied to the frequency coherence of the driving field, which in the single-photon case is directly derived from the laser. Such qubits are also generally susceptible to magnetic field noise, although much of this can be mitigated with modern shielding techniques¹⁰². Readout is achieved in these systems by the electron-shelving technique, driving a transition from one of the two qubit states to a short-lived level with a detectable fluorescent decay^{87,106}.

The hard physical limit on qubit lifetimes can be avoided by encoding the states in separate Zeeman^{52,95} or, if the nucleus has a net spin, hyperfine levels of the same manifold^{83,125}. With transition frequencies ranging from a few megahertz for Zeeman qubits to a few gigahertz for hyperfine qubits, the probability of spontaneous decay is effectively zero and lifetimes can exceed an hour^{130,131}. Due to their nature, Zeeman qubits suffer from the same sensitivity to magnetic field fluctuations as optical qubits¹⁰². With hyperfine structure it is possible to use *clock* states whose transitions are first-order magnetic-field insensitive^{48,90}. In practice, groups generally operate with a small well controlled bias field in order to improve cooling and readout^{30,93}. These qubits can be driven directly with microwaves^{49,81}, but it is more common to use a laser-based Raman transition for reduced crosstalk and stronger coupling to the motion. An alternative approach to reduce crosstalk errors with microwaves is to apply a continuous dressing field to stabilise several field-sensitive states into a two-dimensional basis which is robust against magnetic fluctuations^{124,133}. Such schemes do generally struggle to engineer sufficient ion-motion interaction to perform two-qubit gates quickly, however¹³⁴.

During my PhD, the Imperial experimental ion-trapping group used the states $|g\rangle = 4^2S_{\frac{1}{2}, m_j=-\frac{1}{2}}$ and $|e\rangle = 3^2D_{\frac{5}{2}, m_j=-\frac{1}{2}}$ in $^{40}\text{Ca}^+$ as an optical qubit, with a transition wavelength of 729 nm. Its energy-level structure is illustrated in fig. 3.2. Full details and justifications may be found in recent experimental PhD theses from the group^{57,60,63}, although the original decision was made several PhD cycles ago. $^{40}\text{Ca}^+$ is still an attractive ion for optical qubits, as its transition wavelengths are all between infrared and visible blue, a region well served by commercial optics and lasers¹⁰⁸. This was the setting for the first non-adiabatic two-ion quantum logic gate¹⁰⁹, and the first direct qubit measurements at fidelities high enough to achieve fault-tolerant computing⁸⁶. It continued to be used as a host for investigations

FIGURE 3.2—Energy-level structure of $^{40}\text{Ca}^+$ including the Zeeman sub-levels, with the optical-qubit and readout transitions marked. In the nuclear ground state, it has no nuclear spin and consequently no hyperfine structure. The qubit is encoded in the electric-dipole-forbidden transition at 729 nm, addressed by a single laser. The fluorescent transition at 397 nm is used to measure the qubit; the $|g\rangle$ state interacts radiatively, allowing detection by photon collection, while the $|e\rangle$ state is non-interacting. It is also possible to use the two Zeeman sublevels of the $S_{1/2}$ state as a qubit, with two driving fields in the Raman configuration.



into process tomography⁹⁸, high-fidelity gates⁸ and large-scale entanglement⁸⁵. More recently, a full rack-integrated quantum computing system with two dozen qubits was demonstrated for these calcium ions⁹⁴. The quest for higher fidelities and longer lifetimes, however, has led some other calcium-using groups to move to $^{43}\text{Ca}^+$ due to the hyperfine levels it affords^{49,107}.

For all simple two-level systems, the static contribution to the Hamiltonian is

$$\hat{\mathcal{H}}_{\text{qubit}}/\hbar = \frac{1}{2}\omega_{eg}(|e\rangle\langle e| - |g\rangle\langle g|) = \frac{1}{2}\omega_{eg}\hat{\sigma}_z, \quad (3.1)$$

where the zero point of energy is chosen to be half way between the lower and upper states, respectively labelled $|g\rangle$ and $|e\rangle$. The frequency separation is then ω_{eg} , and it is convenient to write the Hamiltonian in terms of the Pauli Z operator.

3.2 Trapped-ion dynamics

Ion traps confine their contents in three-dimensional space. While the axial fields are almost always supplied electrostatically, the radial confining fields have different properties between Penning and linear rf traps; the former are achieved with static magnetic fields and the latter with oscillating electric fields. A full treatment of ion dynamics should consider these, and would result in radial motion of the ions in addition to the axial motion we will derive in this section. Quantum computing applications almost universally choose to trap ions in a nearly one-dimensional chain along the trap axis, since this allows individual ions to be addressed without complicated stroboscopic techniques. This axis is conventionally labelled z .

We will make the assumption that the radially confining fields are designed such

that the ions do not significantly deviate from the trap axis. For the general positions of any two ions \mathbf{r}_1 and \mathbf{r}_2 , both must have approximately zero radial distance from the axis, and the separation $\mathbf{r}_2 - \mathbf{r}_1 = \delta\mathbf{r}_z + \delta\mathbf{r}_\perp$ should satisfy $|\delta\mathbf{r}_\perp| \ll |\delta\mathbf{r}_z|$ at all times. This assumption allows us to simplify the mathematics down to a single dimension, and represent the usual regime for trapped-ion quantum computing. It breaks down for confining fields that are weak compared to the axial motion, for strong axial confining fields, or for large numbers of ions, but the situations in this thesis will not approach such scenarios. Treatments of the radial motion of trapped ions may be found in theses from the Imperial group on Penning traps⁵⁷ and linear rf traps²⁸.

Under these assumptions, the trap geometry and electromagnetic fields produce a one-dimensional harmonic potential for each ion characterised by a frequency ω_z , which depends on the masses of all of the trapped ions. Taking each ion to be at an axial position z_i , with mass of m_i and a single charge of $-e$, the approximated system potential is

$$V(t) \approx \frac{1}{2} \sum_i m_i \omega_z^2 z_i(t)^2 + \frac{e^2}{4\pi\epsilon_0} \sum_{\langle i,j \rangle} \frac{1}{|z_j(t) - z_i(t)|}. \quad (3.2)$$

This thesis deals only with chains of equal ions, so for simplicity we consider equal masses $m_i = m$. It is further convenient to define a length scale $\ell = \sqrt[3]{e^2/(4\pi\epsilon_0 m \omega_z^2)}$ in order to move to a dimensionless coordinate system defined by $\zeta_i = z_i/\ell$, with the labels ordered such that $\zeta_i < \zeta_{i+1}$. This length scale is generally on the order of micrometres: for $^{40}\text{Ca}^+$ trapped at 500 kHz, it is approximately 7 μm .

With sufficiently strong trapping potentials, the ions will have well-separated equilibrium positions $\zeta_{0,i}$ and undergo small-amplitude oscillations $\delta\zeta_i(t)$ around these points. The equilibrium positions are at the point of zero force, defined by the solutions to

$$\frac{1}{m\omega_z^2\ell^2} \cdot \frac{\partial V}{\partial \zeta_i} = \zeta_i - \sum_{j<i} \frac{1}{(\zeta_j - \zeta_i)^2} + \sum_{j>i} \frac{1}{(\zeta_j - \zeta_i)^2} = 0. \quad (3.3)$$

This must be solved numerically beyond the three-ion case, which is simple with Newton–Raphson iteration. The Jacobian J of the system of equations is

$$J_{ij} = \frac{1}{m\omega_z^2\ell^2} \cdot \frac{\partial^2 V}{\partial \zeta_i \partial \zeta_j} = \begin{cases} 1 + \sum_{k \neq i} \frac{2}{|\zeta_k - \zeta_i|^3} & \text{for } j = i \\ \frac{-2}{|\zeta_j - \zeta_i|^3} & \text{otherwise,} \end{cases} \quad (3.4)$$

Ions	Equilibrium positions $/\ell$									
2				-0.630		0.630				
3			-1.08		0		1.08			
4		-1.44		-0.454		0.454		1.44		
5		-1.74		-0.822		0		0.822		1.74
6		-2.01		-1.14		-0.370		0.370		1.14
7		-2.25		-1.41		-0.687		0		0.687
8		-2.48		-1.66		-0.967		-0.318		0.318
9		-2.68		-1.89		-1.22		-0.600		0
10		-2.87		-2.10		-1.45		-0.854		-0.282

TABLE 3.1—Equilibrium positions of identical singly charged ions in linear chains within a trap. Ions closer to the centre of the chain are closer to their neighbours as the Coulomb force compresses the chain. The positions are given in terms of the length scale $\ell = \sqrt[3]{e^2/(4\pi\epsilon_0 m\omega_z^2)}$. For $^{40}\text{Ca}^+$ trapped at an axial frequency of 500 kHz, $\ell \approx 7 \mu\text{m}$.

and the dimensionality can be halved via symmetry; the positions will be symmetric around the trap centre $\zeta = 0$, with an ion exactly at the centre if there are an odd number. The solutions for up to ten ions are illustrated in table 3.1.

The equations of motion for the small displacements can then be derived from eq. (3.2). Consider the second-order Taylor expansion \bar{V} around the equilibrium, with the reference chosen to make zero potential at zero displacement:

$$\bar{V} \propto \sum_i \left(1 + \sum_{j \neq i} \frac{2}{|\zeta_{0,j} - \zeta_{0,i}|^3} \right) (\delta\zeta_i)^2 - \sum_{\substack{i,j \\ j \neq i}} \frac{2}{|\zeta_{0,j} - \zeta_{0,i}|^3} \delta\zeta_i \delta\zeta_j. \quad (3.5)$$

The dimensionless matrix $(\bar{V}_{ij})/(m\omega_z^2\ell^2)$ is real symmetric over independent displacements and positive semi-definite, so its eigenvalues κ_i^2 are non-negative and its eigenvectors $\{\mathbf{b}_i\}$ form an orthonormal basis. A mapping $\delta\boldsymbol{\zeta}(t) \rightarrow \sum_i q_i(t)\mathbf{b}_i$ then simplifies the classical Hamiltonian to

$$H(t) = \frac{1}{2m} \sum_i p_i(t)^2 + \frac{1}{2} m\omega_z^2 \sum_i \kappa_i^2 q_i(t)^2. \quad (3.6)$$

Ions	Frequency / ω_z	Normal mode participation			
2	1	0.707	0.707	0.707	0.707
	$\sqrt{3}$	-0.707	0.707	0.707	0.707
3	1	0.577	0.577	0.577	0.577
	$\sqrt{3}$	-0.707	0	0.707	0.707
	2.41	0.408	-0.816	0.408	0.408
4	1	0.500	0.500	0.500	0.500
	$\sqrt{3}$	-0.674	-0.213	0.213	0.674
	2.41	0.500	-0.500	-0.500	0.500
	3.05	0.213	-0.674	0.674	-0.213

TABLE 3.2—Joint normal motional modes of identical ions in linear chains within a trap. The frequency of each normal mode is given in terms of the axial trapping frequency ω_z . The participation of each ion in the motion is scaled such that each mode is described by a vector with unit magnitude. The average displacement of each ion is zero; the motion oscillates forwards and backwards.

The standard treatment of quantisation and introduction of creation $\hat{a}_j^\dagger \propto \hat{q}_j + i\hat{p}_j$ and annihilation \hat{a}_j operators diagonalises the motional Hamiltonian to

$$\hat{\mathcal{H}}_{\text{mot}}/\hbar = \sum_i \kappa_i \omega_z \hat{a}_i^\dagger \hat{a}_i \quad (3.7)$$

up to a constant offset, where the sum is over the normal modes of motion. These modes are independent up to the validity of the second-order Taylor expansion; the first neglected terms are on the order of 10^{-3} times weaker.

The ions participate differently in each mode, proportional to the overlap of the relevant individual displacement basis vector and the eigenvector \mathbf{b}_i . The relative frequencies κ_i and the normal-mode participations are depicted in table 3.2. No matter how many ions are in the trap, the two principal modes are the *centre-of-mass* mode at a frequency of ω_z with the ions moving together in phase, and the *breathing* mode at $\sqrt{3}\omega_z$ with the ions expanding and contracting around the centre by amounts proportional to their equilibrium displacement. Higher-energy modes need to be calculated numerically, and the unequal participation amongst ions

typically makes them undesirable for interactions. The later work in this thesis will almost universally deal with the centre-of-mass mode, although with only two ions in a trap, the analyses would apply in the same way on the breathing mode. Importantly, direct measurement of the motional states is not possible. The positional spread of the wavefunction for the zero-point motional state is $q_{\text{RMS}} = \sqrt{\hbar/(2m\kappa_m\omega_z)}$. For $^{40}\text{Ca}^+$ at 500 kHz on the centre-of-mass mode, this is around 16 nm—significantly shorter than the wavelength of any interrogating laser. The only common measurement used in trapped ions is of the qubit states, via the internal states of the ions.

3.3 Ion–laser interactions

3.3.1 General Hamiltonian

A complete description of ion–laser interactions is rather involved, and is left to better works^{17,73,138}. The derivations presented here are illustrative approximations rather than an attempt to be entirely rigorous. For the methods of trapped-ion quantum computing considered in this thesis, only the dynamics of the ion are important; any induced variation in an applied electromagnetic field is relevant only if it has a measurable effect on the ion. The lasers or microwave sources used invariably have sufficient intensity to provide a continuous source of photons, and qubit transitions are deliberately chosen to make spontaneous emission negligible. We will use a semiclassical treatment making use of these features.

In the presence of an electromagnetic field with vector magnetic potential \mathbf{A} the effective momentum of an ion is modified as $\mathbf{p} \rightarrow \mathbf{p} - e\mathbf{A}$, changing the Hamiltonian term to

$$H = \frac{1}{2m}(\mathbf{p} - e\mathbf{A})^2 = \frac{1}{2m}(\mathbf{p}^2 - e\mathbf{p} \cdot \mathbf{A} - e\mathbf{A} \cdot \mathbf{p} + e^2\mathbf{A}^2). \quad (3.8)$$

The \mathbf{p}^2 term is the full dimension of the momentum already considered in eq. (3.6), while the \mathbf{A}^2 term requires at least two-photon processes and is negligible for the desired driving fields. In the Coulomb gauge where \mathbf{A} is purely transverse, it commutes with the quantised momentum $\mathbf{p} \rightarrow \hat{\mathbf{p}} = -i\hbar\nabla$, as $\nabla \cdot (\mathbf{A}\psi) = (\nabla \cdot \mathbf{A})\psi + \mathbf{A} \cdot \nabla\psi$ and the first term is chosen to be zero. This leads to the new ion–field interaction Hamiltonian term being described by

$$\hat{\mathcal{H}}_{\text{int}} = -\frac{e}{m}\hat{\mathbf{p}} \cdot \mathbf{A}. \quad (3.9)$$

As an electromagnetic field at a frequency ω_ℓ , \mathbf{A} has a familiar plane-wave solution

$$\mathbf{A}(\mathbf{r}, t) = \mathbf{A}_0 \exp[i(\mathbf{k} \cdot \mathbf{r} - \omega_\ell t)] + \text{H.c.} \quad (3.10)$$

for a wavevector \mathbf{k} satisfying the dispersion relation $\omega_\ell = c|\mathbf{k}|$. The constant vector \mathbf{A}_0 defines the polarisation axis and strength of the incident field.

An electron's position can be decomposed as its ion's position \mathbf{r}_j plus a small relative displacement $\delta\mathbf{r}_j$. The interaction Hamiltonian for a field acting on multiple ions can then be rewritten as

$$\hat{\mathcal{H}}_{\text{int}} = -\frac{e}{m} \sum_j \underbrace{e^{i(\mathbf{k} \cdot \mathbf{r}_j - \omega_\ell t)}}_{\text{motion}} \underbrace{\hat{\mathbf{p}} \cdot \mathbf{A}_0 e^{i(\mathbf{k} \cdot \delta\mathbf{r}_j - \phi_j)}}_{\text{ion state}} + \text{H.c.}, \quad (3.11)$$

where the position-dependent relative field phase ϕ_j can be absorbed without loss of generality into the chosen basis states. The term labelled *ion state* in this Hamiltonian is responsible for the electronic transitions. Laser light has a wavelength in the hundreds of nanometres, while the distance of the electron from the ion $|\delta\mathbf{r}_j|$ is several orders of magnitude smaller, at around a few Bohr radii. With this exponent, only the first few terms of the Taylor expansion with respect to $\mathbf{k} \cdot \delta\mathbf{r}_j$ are relevant. The leading-order term represents electric dipole transitions, while the term linear in $\mathbf{k} \cdot \delta\mathbf{r}_j$ contains the magnetic dipole and electric quadrupole transitions, and so on. Rather than dealing with the field form explicitly, it is far more convenient to limit the analysis to the two states of interest per ion, $|g\rangle$ and $|e\rangle$, and use Pauli operators to describe the transition. Explicitly, one can define a real coupling strength Ω_j called the *Rabi frequency* such that

$$-\frac{e}{m} \hat{\mathbf{p}} \cdot \mathbf{A}_0 e^{i(\mathbf{k} \cdot \delta\mathbf{r}_j - \phi_j)} \approx \frac{\hbar\Omega_j}{2} e^{-i\phi} \hat{\sigma}_x^{(j)} \quad (3.12)$$

where the ion state definitions are chosen to use the Pauli X operator with a phase difference of ϕ that is the same for all ions. When the states considered are separated by an electric-dipole-forbidden transition, there are also terms proportional to $|g\rangle\langle g|$ and $|e\rangle\langle e|$, which must be compensated experimentally⁴³.

The link to the joint motion in eq. (3.11) is made clearer by considering only the axial component by taking $\mathbf{k} \cdot \mathbf{r}_j \rightarrow k_z \hat{q}_j$, at which point the axial positions \hat{q}_j can be expanded in terms of the raising and lowering operators of each of the motional modes. This leads to a term of the form

$$\exp[i(\mathbf{k} \cdot \mathbf{r}_j - \omega_\ell t)] \rightarrow \exp[i(\sum_m \eta_{j,m} (\hat{a}_m + \hat{a}_m^\dagger) - \omega_\ell t)], \quad (3.13)$$

where the sum is over the motional modes, and

$$\eta_{j,m} = k_z s_{j,m} \sqrt{\frac{\hbar}{2m\omega_z}} \quad (3.14)$$

is the *Lamb–Dicke parameter*, which characterises the coupling of each ion-state transition with each motional-mode transition. Ion traps typically operate with its value around $1/10$ ^{28,107}. The dimensionless scaling factors $s_{j,m}$ are defined by how strongly ion j couples to motional mode m . Using the relative mode frequencies κ_m and the normal-mode participation unit vector components $b_{j,m}$ given in table 3.2, the scaling factors for a chain of N ions⁵⁹ are $s_{j,m} = b_{j,m} \sqrt{N/\kappa_m}$. As expected, the centre-of-mass mode is most strongly affected by the incident laser field due to having the smallest $\kappa = 1$. For simplicity, we will consider only this mode from now on, for which all ions have the same Lamb–Dicke parameter η . Combining the three Hamiltonian components of eqs. (3.1), (3.7) and (3.11), we reach the lab-frame ion-trap Hamiltonian for two-level ions and their centre-of-mass motion:

$$\hat{\mathcal{H}}_{\text{lab}}/\hbar = \underbrace{\frac{1}{2} \sum_j \omega_{eg} \hat{\sigma}_z^{(j)} + \omega_z \hat{a}^\dagger \hat{a}}_{\hat{\mathcal{H}}_{\text{sys}}/\hbar} + \underbrace{\sum_j \Omega \hat{\sigma}_x^{(j)} \cos[\eta(\hat{a} + \hat{a}^\dagger) - \omega_\ell t - \phi]}_{\hat{\mathcal{H}}_{\text{laser}}/\hbar}. \quad (3.15)$$

The basis of interest is the eigenstates of the system Hamiltonian. These are labelled $|x, n\rangle$, where n is the number of phonons in the centre-of-mass mode, and $x \in \{g, e\}^N$ is a descriptor of the state of the N ions. For example, $|gg, 0\rangle$ is the joint ground state of two ions and the motion.

In order to more clearly calculate the allowed transition processes, the electronic and motional energy terms can be removed from the Hamiltonian by moving to an interaction frame. Under a unitary transformation $\exp(i\hat{\mathcal{H}}_{\text{sys}}t/\hbar)$, the new interaction Hamiltonian becomes

$$\begin{aligned} \hat{\mathcal{H}}_{\text{int}}/\hbar = \Omega \sum_j e^{i\omega_{eg}t\hat{\sigma}_z^{(j)}/2} \hat{\sigma}_x^{(j)} e^{-i\omega_{eg}t\hat{\sigma}_z^{(j)}/2} \\ \times e^{i\omega_z t \hat{a}^\dagger \hat{a}} \cos[\eta(\hat{a} + \hat{a}^\dagger) - \omega_\ell t - \phi] e^{-i\omega_z t \hat{a}^\dagger \hat{a}} \end{aligned} \quad (3.16)$$

once all trivial commutations have been resolved. Recognising that $\hat{\sigma}_z^2 = 1$ and writing $\hat{\sigma}_x = \hat{\sigma}_+ + \hat{\sigma}_-$, a series expansion gives $\exp(i\chi\hat{\sigma}_z) = \cos\chi + i\hat{\sigma}_z \sin\chi$, so

$$\begin{aligned} e^{i\chi\hat{\sigma}_z} \hat{\sigma}_x e^{-i\chi\hat{\sigma}_z} &= \cos(2\chi)(\hat{\sigma}_+ + \hat{\sigma}_-) + i\sin(2\chi)(\hat{\sigma}_+ - \hat{\sigma}_-) \\ &= e^{2i\chi} \hat{\sigma}_+ + e^{-2i\chi} \hat{\sigma}_-. \end{aligned} \quad (3.17)$$

The motional term requires more complex machinery to evaluate. Using the Baker–Campbell–Hausdorff-derived eq. (2.13) to expand $e^{\hat{A}}\hat{B}e^{-\hat{A}}$ into commutators of the two operators, and as $[\hat{a}^\dagger\hat{a}, \hat{a}] = -\hat{a}$ and $[\hat{a}^\dagger\hat{a}, \hat{a}^\dagger] = \hat{a}^\dagger$, we have

$$e^{i\chi\hat{a}^\dagger\hat{a}}(\hat{a} + \hat{a}^\dagger)e^{-i\chi\hat{a}^\dagger\hat{a}} = e^{-i\chi}\hat{a} + e^{i\chi}\hat{a}^\dagger. \quad (3.18)$$

We are free to insert the identity expression $e^{-i\chi\hat{a}^\dagger\hat{a}}e^{i\chi\hat{a}^\dagger\hat{a}}$ between all operators of a power series, so the cosine term in the lab-frame Hamiltonian of eq. (3.15) is moved to the interaction picture by the replacement $\hat{a} \rightarrow e^{-i\omega_z t}\hat{a}$.

There are now three relevant frequencies: the qubit-state separation ω_{eg} , the motional-state separation ω_z , and the interaction field ω_ℓ . While there are many possible transitions, only those close to resonance can meaningfully contribute to the dynamics. To simplify the interaction-picture Hamiltonian eq. (3.16), we make a rotating-wave approximation to neglect any terms with $\omega_{eg} + \omega_\ell$, and define a new selection frequency $\omega_s = \omega_\ell - \omega_{eg}$. This leads to a final Hamiltonian

$$\hat{\mathcal{H}}_{\text{int}}/\hbar = \sum_j \frac{\Omega}{2} e^{-i(\omega_s t + \phi)} \hat{\sigma}_+^{(j)} \exp\left[i\eta(e^{-i\omega_z t}\hat{a} + e^{i\omega_z t}\hat{a}^\dagger)\right] + \text{H.c.}, \quad (3.19)$$

where the sum is over ions targeted by the interaction. This is the base Hamiltonian for trapped-ion quantum computing. The derivations here used a single interaction field, but essentially the same Hamiltonian is reached if the field is formed by two separate components in a Raman configuration. The relevant wavevector becomes the difference between the two fields, but the transitions remain in the same form.

3.3.2 Sideband transitions

The general interaction Hamiltonian of eq. (3.19) still contains two frequencies and several possible transitions. For an initial qualitative view, the exponential can be expanded up to the term linear in the Lamb–Dicke parameter. This gives three terms in a rotating-wave approximation, shown here for a single ion:

$$\hat{\mathcal{H}}/\hbar \approx \begin{cases} i\eta\frac{\Omega}{2}(e^{-i(\delta t + \phi)}\hat{\sigma}_+\hat{a} - e^{i(\delta t + \phi)}\hat{\sigma}_-\hat{a}^\dagger) & \omega_s = -\omega_z + \delta, \text{ red;} \\ \frac{\Omega}{2}(e^{-i(\delta t + \phi)}\hat{\sigma}_+ + e^{i(\delta t + \phi)}\hat{\sigma}_-) & \omega_s = \delta, \text{ carrier;} \\ i\eta\frac{\Omega}{2}(e^{-i(\delta t + \phi)}\hat{\sigma}_+\hat{a}^\dagger - e^{i(\delta t + \phi)}\hat{\sigma}_-\hat{a}) & \omega_s = \omega_z + \delta, \text{ blue.} \end{cases} \quad (3.20)$$

The *carrier* transition is driven if the interaction-field frequency is close to the qubit frequency, simply driving coherent oscillations between the two qubit states without affecting the motion. If the interaction field is instead tuned to be one




FIGURE 3.3—Sideband transitions up to first order in a single trapped ion inside the Lamb–Dicke regime. The carrier couples the electronic states without affecting the motion, and is driven when the interaction-field frequency is close to the separation between these two states. When the driving frequency is detuned by one motional quantum from the qubit frequency, the red or blue sidebands can be driven, which respectively remove or add a phonon while exciting the ion.

motional frequency away from the qubit frequency, one of the two first-order sidebands is driven instead. These are called the *red* and the *blue*, with the red sideband having the lower frequency. In these, a phonon of motion is removed from or added to the system as the ion is excited, respectively. These low-order terms are illustrated in fig. 3.3.

In each of these transitions, δ refers to some separate small detuning, with the validity of this second rotating-wave approximation being approximately defined by $\delta \ll \omega_z$. Strictly, the approximation is valid only when every sideband is well separated in frequency from every other transition. As more motional modes are considered, especially at higher-order sidebands than given here, the frequency spectrum becomes denser and off-resonant terms can play a more significant role. In this thesis, we will only need to work in regimes where this approximation is valid, either because of a weak interaction strength, or by exciting the transitions very close to resonance. Each transition is on resonance when its δ is zero, though mis-calibrations and other drifts often mean that the transitions are not addressed precisely resonantly.

Each of these transitions couples pairs of joint qubit–motion states in a simple Rabi model. For example, the carrier drives the transitions $|g, n\rangle \leftrightarrow |e, n\rangle$, and the blue sideband drives $|g, n\rangle \leftrightarrow |e, n+1\rangle$. Taking a single coupled pair of states and reducing the labels to $|g\rangle$ and $|e\rangle$, the Schrödinger equation can be solved exactly

as a simple pair of coupled differential equations, giving a time-evolution operator

$$\begin{aligned}\hat{\mathcal{U}}(t) = & e^{-i\delta t/2} \left(\cos(\Omega' t/2) + i \frac{\delta}{\Omega'} \sin(\Omega' t/2) \right) |e\rangle\langle e| \\ & + e^{i\delta t/2} \left(\cos(\Omega' t/2) - i \frac{\delta}{\Omega'} \sin(\Omega' t/2) \right) |g\rangle\langle g| \\ & - i \frac{\Omega}{\Omega'} \sin(\Omega' t/2) \left(e^{-i\delta t/2} e^{-i\phi'} |e\rangle\langle g| + e^{i\delta t/2} e^{i\phi'} |g\rangle\langle e| \right).\end{aligned}\quad (3.21)$$

Exactly on resonance, this describes perfect sinusoidal oscillations between the two coupled states with a *Rabi frequency* of Ω . If the laser is detuned from the transition by an amount δ , the Rabi frequency is modified to $\Omega' = \sqrt{\delta^2 + \Omega^2}$, and the oscillation amplitude is reduced by a factor of $1 + \delta^2/\Omega^2$.

Conventional terminology is to call the shortest pulse that completely exchanges the state populations a π *pulse*, although from a mathematical point of view, its duration is $t = \pi/\Omega'$ implying $\hat{\mathcal{U}}$ is 4π -periodic. After twice the length of a π pulse, the measured populations are the same, but a global phase factor of -1 is introduced on the ion state. This is most relevant when the transition is applied to a single ion in a chain.

Each pair of ground and excited states is coupled by exactly one transition, and the oscillation frequency $\Omega_{n,m}$ is dependent only on the two motional levels n and m and the Lamb–Dicke parameter. To determine these more accurately, we consider a more complete expansion of the exponential in the interaction Hamiltonian of eq. (3.19). Using the Baker–Campbell–Hausdorff formula with $[[\hat{a}, \hat{a}^\dagger], \hat{a}] = 0$, the motional component can be written as

$$\exp\left[i\eta(e^{-i\omega_z t} \hat{a} + e^{i\omega_z t} \hat{a}^\dagger)\right] = \exp(-\eta^2/2) \exp(i\eta e^{i\omega_z t} \hat{a}^\dagger) \exp(i\eta e^{-i\omega_z t} \hat{a}). \quad (3.22)$$

The new coupling frequency is $\Omega_{n,m} = \Omega |\langle m| \circ |n\rangle|$, where the matrix element is of eq. (3.22). In this form, it is clear that for any pair of starting n and ending m motional states, all transition processes contain only frequencies $(m - n)\omega_z$, corresponding to the phonon difference. Further, there are a finite number of contributing motion-dependent terms as only the first n terms in the series of expansion of $\exp(\chi \hat{a})|n\rangle$ have a nonzero coefficient.

Explicitly, the frequencies are

$$\frac{\Omega_{n,m}}{\Omega} = e^{-\eta^2/2} \eta^{|m-n|} \sqrt{\frac{\min(n,m)!}{\max(n,m)!}} L_{\min(n,m)}^{(|m-n|)}(\eta^2), \quad (3.23)$$

when defined in terms of the generalised Laguerre polynomials

$$L_n^{(a)}(x) = \sum_{j=0}^n (-1)^j \binom{n+a}{n-j} \frac{x^j}{j!}. \quad (3.24)$$

The Rabi frequency describes a single physical coupling between two levels, and the symmetry of eq. (3.23) mathematically illustrates that $\Omega_{n,m} = \Omega_{m,n}$. The first-order transition frequencies are


$$\begin{aligned} \text{carrier:} \quad & \frac{\Omega_{n,n}}{\Omega} = 1 - \frac{1}{2}(2n+1)\eta^2 + \mathcal{O}(\eta^4), \\ \text{red and blue:} \quad & \frac{\Omega_{n,n+1}}{\Omega} = \eta\sqrt{n+1} \left[1 - \frac{1}{2}(n+1)\eta^2 + \mathcal{O}(\eta^4) \right]. \end{aligned} \quad (3.25)$$

The leading-order terms of the second-order sideband transition frequencies are $\eta^2\sqrt{(n+1)(n+2)}$. The suitability of the series-expansion approximation used to define eq. (3.20) therefore depends on both η and the number of phonons being small. Its region of validity is named the *Lamb–Dicke regime*. There are many mathematical definitions of this in the literature, but the requirements are usually that second-order sideband transitions are forbidden, and the frequencies of the carrier and the first-order sidebands can be truncated to their leading-order terms. One simple expression of these is $(n+1)\eta^2 \ll 1$. Within this limit, the carrier couples all its pairs of states at the same frequency, but the first-order sidebands couple proportional to the square root of the larger number of phonons, so different pairs have generally incommensurate oscillations.

3.4 Mølmer–Sørensen gate

Only the carrier transition affects every possible basis state of the ion trap. Notably, the first red sideband does not affect the state $|g, 0\rangle$, while the first blue sideband does not affect $|e, 0\rangle$. The joint motion can therefore be used as a communication bus, allowing two-qubit gates to be realised with only single-ion operations by entangling the internal states of the ions with the motion. This was the earliest major proposal for fast, scalable quantum computing²⁵ and became the first implemented ion-trap quantum-logic gate, although the initial demonstration was between an ionic and a motional qubit⁸³. More advanced laser stabilisation was needed before two separate ions could be coherently addressed and entangled with this scheme¹⁰⁹. Further, for general computing use, the reliance on a coherent motional qubit is undesirable. Motional states decohere quickly due to voltage fluctuations and trap

FIGURE 3.4—The energy levels of the Mølmer-Sørensen scheme. Two global fields are applied: one slightly detuned from the blue sideband, and the other detuned by an equal but opposite amount from the red sideband. Inside the Lamb-Dicke regime there are oscillatory dynamics between $|gg, n\rangle \leftrightarrow |ee, n\rangle$ and $|ge, n\rangle \leftrightarrow |eg, n\rangle$ that are independent of the motional state $|n\rangle$. If the interaction strength and detuning are chosen appropriately, there is a pulse duration that will produce two-qubit entanglement with no spurious coupling to the motion.



heating, while the requirement to begin in the ground state of motion imposes onerous cooling requirements. A better scheme would use the motion to couple the qubits without the interaction strength being conditioned on the motional state.

One such option is the *Mølmer-Sørensen gate*^{115,116}. Assuming the Lamb-Dicke regime, one applies both the red and blue sidebands simultaneously to multiple ions, detuned by equal but opposite amounts ϵ at the same phase. The scheme is illustrated in fig. 3.4. This produces a Hamiltonian

$$\hat{\mathcal{H}}_{\text{MS}}/\hbar = \frac{\eta\Omega}{2} (ie^{-i\phi}\hat{S}_+ - ie^{i\phi}\hat{S}_-) (e^{-i\epsilon t}\hat{a}^\dagger + e^{i\epsilon t}\hat{a}), \quad (3.26)$$

where $\hat{S}_\circ = \sum_j \hat{\sigma}_\circ^{(j)}$ is a sum of single-qubit Pauli operators. The phase ϕ chooses a qubit operator in the $\hat{\sigma}_x$ - $\hat{\sigma}_y$ plane, and the Mølmer-Sørensen gate is accordingly occasionally referred to as a $\hat{\sigma}_\phi \otimes \hat{\sigma}_\phi$ gate to distinguish it from $\hat{\sigma}_z \otimes \hat{\sigma}_z$ -interaction schemes¹⁰¹. We will arbitrarily choose ϕ to make the qubit terms \hat{S}_y .

The time evolution of this Hamiltonian can be found by the Magnus expansion, which terminates after two terms. With $\hat{\mathcal{U}}_{\text{MS}}(t) = \exp(\hat{M}_1(t) + \hat{M}_2(t))$, the Magnus operators from eq. (2.16) are

$$\begin{aligned} \hat{M}_1(t) &= -i\frac{\eta\Omega}{2}\hat{S}_y \int_0^t dt_1 (e^{-i\epsilon t_1}\hat{a}^\dagger + e^{i\epsilon t_1}\hat{a}), \text{ and} \\ \hat{M}_2(t) &= i\frac{(\eta\Omega)^2}{4}\hat{S}_y^2 \int_0^t dt_1 \int_0^{t_1} dt_2 \sin(\epsilon(t_2 - t_1)). \end{aligned} \quad (3.27)$$

The first of these describes a state-dependent phase-space displacement: the positive and negative eigenstates of \hat{S}_y undergo opposite circular trajectories. The second term provides two-qubit interactions via the $\hat{S}_y^2 = 2(1 + \hat{\sigma}_y \otimes \hat{\sigma}_y)$ operator. If the interaction is applied for a time $t = 2\pi/\epsilon$, the phase-space displacement returns to

zero and the whole evolution is

$$\hat{\mathcal{U}}_{\text{MS}}\left(\frac{2\pi}{\epsilon}\right) \equiv \cos\left(\pi \frac{\eta^2 \Omega^2}{\epsilon^2}\right) - i \sin\left(\pi \frac{\eta^2 \Omega^2}{\epsilon^2}\right) \hat{\sigma}_y \otimes \hat{\sigma}_y, \quad (3.28)$$

up to a global phase. Geometrically, the angle inside the trigonometric functions is proportional to the area swept out by the phase-space trajectory. If the detuning is set to $\epsilon = 2\eta\Omega$, the Mølmer–Sørensen interaction creates Bell states from unentangled ions. With simple additional single-qubit rotations, this can be transformed into a quantum logic gate. Importantly, assuming the Lamb–Dicke regime holds, this interaction is not dependent on the motional state, and so has far less taxing requirements on cooling and isolation from external fields. This is not limited to two-qubit processes. Provided all ions partake equally in the motional mode addressed, the same technique maps the joint electronic ground state to a GHZ-type state*, independent of the motional occupation and the number of ions^{82,114}. Practically, the centre-of-mass mode is not always ideal due to its propensity to heating, but for four ions there is a *stretch* mode with equal participation in alternating directions that can be used for larger-scale entanglement¹⁰³.

Aside from requiring relatively weak ion–motion coupling to achieve the Lamb–Dicke regime, the other major requirement for good fidelity is that the field does not significantly drive the carrier transition off-resonantly. This requires that $\Omega \ll \omega_z$. The original formulation of the gate¹¹⁵ drove the transition adiabatically, preventing population of the intermediate states. It relied on a phase difference between the left and right paths of fig. 3.4, which, when treated as two-photon processes, have coupling frequencies proportional to n and $n + 1$ respectively. The phase difference cancelled out the motion-dependent n component of these. This was only possible with a further constraint that $\eta\Omega \ll \epsilon$, and so the gate was exceedingly slow. The stronger-coupled form removes this restriction¹¹⁶, allowing the gate speed to be generally limited by available laser power and the validity of the Lamb–Dicke approximation. At the time of writing, the record fidelity for a two-qubit gate was achieved by this method⁴⁰, tied with a $\hat{\sigma}_z \otimes \hat{\sigma}_z$ -based gate also in trapped ions⁴.

Chapter 5 will examine the dynamics of the Mølmer–Sørensen interaction when there are various static frequency offsets and miscalibrations, and present driving schemes to mitigate undesired effects. In chapter 6, the gate will be taken *outside* the Lamb–Dicke approximation by a general method that produces the same Bell-state

*Greenberger–Horne–Zeilinger states are of the form $(|00 \dots 0\rangle + |11 \dots 1\rangle)/\sqrt{2}$, and the term is usually only applied to three-or-more-qubit systems. The Mølmer–Sørensen interaction creates states of the form $(|00 \dots 0\rangle + e^{i\phi}|11 \dots 1\rangle)/\sqrt{2}$, including the two-qubit case.

creation, breaking previously fundamental limitations on motional populations and usable ion–motion coupling strengths. First, however, we return to single-ion dynamics, and consider the problem of coherence creation and certification.

Part II

New Research

Chapter 4

Certification of Higher-Order Coherence

COAUTHORSHIP

The experimental work in this chapter was carried out by the ion-trapping group at Imperial College London. I calculated all of the measurement statistics to derive the unbiased estimators, and found the optimal state-creation and projection sequences numerically. Florian Mintert and I proved the robustness of the interference-pattern certifier under general measurement operators, and I performed all the numerical optimisations to verify the new threshold values. The work in this chapter was also described in ref. 29.

The manipulation and distillation of entanglement was recognised early as essential to quantum computation¹¹, leading to efforts to move beyond Bell-inequality tests to more discriminating methods of detecting entanglement^{54,92} that have only multiplied with time^{55,56}. Coherence, however, has only more recently been recognised as a resource in the same manner^{6,69,119}, with its uses now being known for applications varying from quantum information processing^{51,111}, to the creation of nonequilibrium entropy¹⁰⁵ and the extraction of thermodynamic work⁶⁶.

A hierarchy for quantum coherence was defined in section 2.1, where a pure state $|\psi\rangle$ is said to be coherent in a particular basis $\{|j\rangle\}$ if its representation $|\psi\rangle = \sum_j c_j |j\rangle$ has at least two nonzero coefficients c_j . This readily extends to multiple levels, where a pure state is called k -coherent, or said to have a coherence rank of k , if it has at least k non-zero coefficients. Mixed states are k -coherent if all possible pure-state decompositions include at least one k -coherent element. Despite having this simple classification scheme, it is not trivial to continuously quantify coherence in the general case. It feels logical that $(|0\rangle + |1\rangle)/\sqrt{2}$ should be somehow more coherent than $\sqrt{0.1}|0\rangle + \sqrt{0.9}|1\rangle$ due to the greater imbalance between the two basis states, yet this idea is harder to justify at higher levels of coherence, or with density operators that may have multiple possible decompositions. Direct measurement is typically fraught, as the only measurement basis available is usually the basis over which the coherence is defined. Still, it is valuable to be able to classify the degree

of coherence of a system; higher-order coherence is its own resource, which may be expended to enhance phase-discrimination tasks^{20,99}.

Assuming perfectly coherent operations and lossless measurements in the coherence basis, it would be possible to determine the coherence of a state by complete reconstruction of its density operator. Even with these perfect conditions, extensive time is required to build up sufficiently precise statistics on the elements, and for high orders of coherence in mixed states the subsequent classical analysis is similarly difficult. These obstacles motivate a different approach, just as they do in entanglement verification. Instead of inferring the state, we can instead take a different measurement whose sole purpose is to distinguish states with different orders of coherence.

These coherence certifiers are entirely analogous to entanglement witnesses, although there are significant further complications in the former. Entanglement may be detected by coherent local operations and classical communications, but its generation requires operations outside this set⁹². Coherence, however, is detected and generated by the same set of operations. This appears to enforce an unfortunate circular requirement that any measurement to verify the preparation of a coherent superposition must trust the same operations it is assessing.

This chapter describes a robust high-order coherence certifier that overcomes the issues both in scalability of measurements and in impossible assumptions of the naïve approach. It works when even the basis of coherence is not accessible to measurement, and is demonstrated by an experimental realisation in the motional state of a trapped ion in collaboration with the group at Imperial. The metric is provably immune to false positives, and requires only simple functions of a one-dimensional interference pattern. It is built on prior work out of the Imperial quantum information theory group³⁴, with significant extra effort required to make the scheme valid for general quantum measurements that cannot even distinguish the coherence basis states.

4.1 Quantum coherence

The concept of coherence itself is not unique to quantum mechanics. Classical coherence is a core component of the wave theory of electromagnetism. Quantum coherence, on the other hand, allows for single particles to interfere with themselves. Along with entanglement and quantisation itself, it plays a central role in quantum effects, from the complete description of the laser to the celebrated Hong–Ou–Mandel dip⁵³.

The quantification of this coherence was first explored by noting that two-level coherence may be characterised by the magnitude of the off-diagonal term in the density operator, and then generalising this across pairs of orthogonal subspaces that spanned the entire Hilbert space¹. While this initial effort was not considered as such, the popularity of resource theories²² led to more complete works to fit coherence into this new framework^{6,69}. The principal additions are the identification of a set of incoherent operations that cannot increase coherence, analogous to the local operations and classical communication with entanglement, and a requirement that a measure is convex under the mixing of states. Both are physically motivated: one cannot conjure coherence from nothing, nor can they increase it by classically mixing two states together. Formally, the incoherent operations are the quantum channels that map the set of incoherent states to itself, and a coherence measure C is convex if

$$\lambda C[\hat{\rho}_1] + (1 - \lambda) C[\hat{\rho}_2] \geq C[\lambda \hat{\rho}_1 + (1 - \lambda) \hat{\rho}_2], \quad (4.1)$$

for a mixing parameter $0 \leq \lambda \leq 1$ and all pairs of density operators $\hat{\rho}_1$ and $\hat{\rho}_2$.

The convexity condition is reminiscent of the triangle inequality, and indeed distance-based measures are perhaps the most studied of coherence quantifiers from a theoretical perspective¹¹⁹. The insight is to quantify the coherence of a state by the minimum distance between it and an incoherent state. This permits a family of measures, based on the particulars of the distance used. Various works have investigated the properties of using the relative entropy^{6,137}, the state infidelity¹²⁰ and both Schatten- p and ℓ_p matrix norms^{6,97}.

Measures and witnesses of higher-order coherence are more complex to construct, as they target an understanding of a more nuanced view of coherence. As an illustration, the ℓ_1 norm is commonly used to characterise two-level coherence, but plainly cannot distinguish the number of superposition elements taken alone. More involved functionals of the off-diagonal terms have been used to construct more complete quantifiers^{69,99}, but these still require detailed tomographic measurements. Some of these schemes may be used to construct witnesses, or to use incomplete data to lower-bound the rank of coherence present⁹⁹, but these prior methods make significant assumptions about the properties of the physical system and the correctness of coherent manipulations. Instead, we turn to interference-pattern methods, which are well-known as a standard indicator of coherence in two-level systems.

4.2 Interference-pattern methods

Ramsey-type experiments have long been a standard tool to verify coherence between two states. In these, the system is first subjected to some coherent operation that ought to create a superposition between the basis states. It evolves freely by a varied amount, then the reverse of the original coherent operation is applied and the population of one of the two states is measured. Regardless of how faithfully the coherent operation was performed, evidence of oscillation in the resulting data is proof that coherence must have existed; an incoherent state would be unable to interfere with itself during the reverse mapping. The simplicity of these experiments and the minimal amount of data required make them very attractive as a base for inference of further system properties⁸⁹. We will consider how they may be used to classify multilevel coherence.

In a two-level Ramsey experiment, the two states are allowed to evolve relative to each other for a complete period. For photonic systems, this is equivalent to creating a path difference between the two basis states and applying a controllable phase shift to only one of the paths before recombination. Physically this describes a Mach-Zehnder interferometer. In higher dimensions, one can imagine a generalisation of such an interferometer as having a different path and phase shift for each basis state¹²⁹. This creates an interference pattern over as many variables as there are dimensions in the Hilbert space. Notably, while any evidence of periodic oscillation in a two-level Ramsey experiment proves some degree of coherence, one cannot distinguish true higher-order coherence from an incoherent mixture of pairwise coherent systems by simply counting the frequency components. The maximal peak-to-peak visibility does, however, encode some information about the rank of the coherence of the state; there are threshold values for each rank of coherence that no lower-ranked coherent state can exceed.

In arbitrarily many dimensions, optimising to find the global maximum becomes experimentally taxing. It is similar in principle to general state tomography methods. To ease the computational burden, one can instead consider only lower-order moments $\{M_n\}$ of an interference pattern, defined by

$$M_n(\hat{\rho}) = \int d\mathbf{w}(\boldsymbol{\phi}) \langle \boldsymbol{\phi} | \hat{\rho} | \boldsymbol{\phi} \rangle^n, \quad \text{with} \quad |\boldsymbol{\phi}\rangle = \frac{1}{\sqrt{d}} \sum_j e^{-i\phi_j} |\phi_j\rangle, \quad (4.2)$$

and some measure $w(\boldsymbol{\phi})$ that acts as a prior. The uniform prior $(2\pi)^{-d}$ is appropriate for completely unknown input states of dimension d , but others can be used to reduce the sampling requirements in the discretisation of the integral when approx-

imating it experimentally. Lower-order moments vary less rapidly with respect to the state, and so require fewer measurements to approximate well. Calculating $\lim_{n \rightarrow \infty} M_n^{1/n}$ is equivalent to finding the maximal peak-to-peak visibility, which has the theoretically best distinguishing characteristics. The lower moments still present the same thresholded structure, but with lower distinguishability¹²⁹. The moments also satisfy the convexity property of eq. (4.1), a critical condition for them to certify coherence in the context of a resource theory. However, retrieving any of this information requires either that projective measurements of arbitrary states can be taken accurately, or that individual basis states can be phase-shifted independently with completely reliable coherent manipulations. These severely limit the viability of this multi-dimensional approach.

These problems were addressed by prior work out of the Imperial theory group³⁴. This considers any space whose coherence basis vectors $\{|j\rangle\}$ can be made to evolve under the equally spaced Hamiltonian $\hat{\mathcal{H}} \propto \sum_j j|j\rangle\langle j|$. If necessary, one can also expand the Hilbert space with intermediate dummy states that are never populated to achieve the equal separation. This restriction on the required evolution replaces the projection onto an arbitrary state $|\phi\rangle$ in eq. (4.2) with a free evolution $\hat{\mathcal{U}}_f(\phi) = \sum_j e^{-ij\phi}|j\rangle\langle j|$ followed by some fixed mapping sequence $\hat{\mathcal{U}}_m$, making the interference pattern one-dimensional. The Hamiltonian naturally occurs as the free evolution in harmonic-oscillator systems, but can be effectively driven in many others, including those with degeneracy. For example, in a system of d qubits and the coherence defined over the product of z -basis eigenstates, the required evolution can be realised by implementing

$$\hat{\mathcal{U}}(\phi) = \hat{\mathcal{R}}_1(\phi)\hat{\mathcal{R}}_2(2\phi) \cdots \hat{\mathcal{R}}_d(2^{d-1}\phi) \quad \text{where} \quad \hat{\mathcal{R}}_k(\phi) = \exp\left(i\phi\hat{\sigma}_z^{(k)}\right). \quad (4.3)$$

The operators $\hat{\mathcal{R}}$ are rotations around the individual qubits' Pauli Z axes, and the whole evolution requires only single-qubit operations. This form of evolution is frequently implemented virtually, in an error-free manner. Control pulses generally evolve synchronously with the basis-state phase evolution, in which case applying a constant phase shift to the driving fields is equivalent to separate evolution.

In this simpler system, the interference pattern for projective measurements onto $|\chi\rangle$ becomes

$$p(\phi) = \langle \chi | \hat{\mathcal{U}}_m \hat{\mathcal{U}}_f(\phi) \hat{\rho} \hat{\mathcal{U}}_f^\dagger(\phi) \hat{\mathcal{U}}_m^\dagger | \chi \rangle. \quad (4.4)$$

With only a single dimension, it is now always feasible to evaluate the entire interference pattern, so the prior from eq. (4.2) can be replaced with the standard

uniform distribution, leaving the moments as

$$M_n = \frac{1}{2\pi} \int_0^{2\pi} d\phi p(\phi)^n. \quad (4.5)$$

These moments alone no longer provide the desired threshold structure; they are all maximised to unity by taking $\hat{\rho} = |\chi\rangle\langle\chi|$, which is completely incoherent. Instead, Dive *et al.*³⁴ showed that the family of ratios of moments

$$R_n = \frac{M_n}{M_1^{n-1}} \quad (4.6)$$

satisfy all the necessary conditions to be certifiers of higher-order coherence: they are convex in both arguments by the definition of eq. (4.1), and they have hierarchical threshold values such that the maximal value of R_n with a k -coherent state is strictly less than the maximal value with a $(k+1)$ -coherent state. All ratios with $n > 2$ are capable of certifying all ranks of coherence. In practice, we will only use R_3 for the same experimental reason that lower moments are generally preferred: they can be well approximated with fewer measurements.

Proving the convexity of the $\{R_n\}$ certifiers over the tested state is largely straightforward. The pattern $p(\hat{\rho}, \phi)$ is always between zero and one for all normalised states, and p^n is trivially convex for non-negative p . The moments M_n are therefore convex as integration with respect to ϕ is linear with respect to the state. Showing the convexity of R_n can then be achieved by showing that the second derivative

$$\partial_\lambda^2 R_n(\lambda \hat{\rho}_1 + (1-\lambda) \hat{\rho}_2) \geq 0 \quad \text{for } 0 \leq \lambda \leq 1, \quad (4.7)$$

since the function is asymptote free. This can be evaluated in terms of the moments—dropping the function arguments for clarity—to give

$$\partial_\lambda^2 R_n = \frac{1}{M_1^{n+1}} \left[M_1^2 (\partial_\lambda^2 M_n) - 2(n-1) M_1 (\partial_\lambda M_1) (\partial_\lambda M_n) + n(n-1) M_n (\partial_\lambda M_1)^2 \right], \quad (4.8)$$

which can be reduced to a trivially positive form

$$= \frac{n(n-1)}{M_1^{n+1}} \left\langle p^{n-2} [p \langle \partial_\lambda p \rangle - \langle p \rangle (\partial_\lambda p)]^2 \right\rangle \quad (4.9)$$

using the notation $\langle f \rangle = \int_0^{2\pi} d\phi f(\phi)/(2\pi)$, the explicit derivatives

$$\partial_\lambda M_n = n \langle p^{n-1} (\partial_\lambda p) \rangle \quad \text{and} \quad \partial_\lambda^2 M_n = n(n-1) \langle p^{n-2} (\partial_\lambda p)^2 \rangle, \quad (4.10)$$

and rearrangements of the form $\langle c \rangle = c$ for interference-phase-independent c .

It is generally intractable to analytically calculate the hierarchical threshold values implied by the convexity of the certifier and the convexity of the set of k -coherent states. The maximum value of any R_n for an incoherent state is trivially one, and Dive *et al.*³⁴ showed that the maximal value of R_3 for 2-coherence is $\frac{5}{4} = 1.25$ and upper-bounded the maximal value for 3-coherent states to $\frac{179}{96} \approx 1.86$. They also performed thorough numerical optimisations—I personally replicated these using the techniques described later—to find empirical upper bounds on the actual attainable maxima for each coherence rank: 1.25, 1.77, 2.32 and 2.88 were the largest observed values for 2-, 3-, 4- and 5-coherent states respectively. It is worth highlighting that R_3 is not a measure of coherence, but a form of witness; it is not necessarily zero for an incoherent state, and there are incoherent states that exhibit larger values of R_3 than coherent states. Our intent is to use it to certify balanced coherent superpositions as being unambiguously of the desired rank, rather than to produce a perfectly discriminating measure.

The convexity of the certifier with respect to the input state and the convexity of the sets of k -coherent states make R_3 a valid witness. However, this is not sufficient to make the certifier valid in cases where the coherent mapping operation $\hat{\mathcal{U}}_m$ cannot be trusted. This is an unfortunate problem for most certifiers, since the mapping is invariably implemented with the same operations that prepare the state to be tested, so if the state cannot be prepared coherently—what this certifier purports to test—then the likelihood of $\hat{\mathcal{U}}_m$ functioning correctly is also high. It is therefore imperative that a faulty mapping to the measurement basis cannot increase the coherence. Dive *et al.*³⁴ showed this for the case that a projective measurement onto the state $\hat{\mathcal{U}}_m|\chi\rangle$ was replaced by a probabilistic projection onto one of a set of states $|\chi_j\rangle$ with probability p_j , with the proof progressing near-identically to the proof of convexity over the tested state. This means that under these relatively relaxed assumptions about the measurement, the R_n certifiers can *never* produce false positives when detecting coherence.

These assumptions do not hold in our experimental realisation of choice, though. We seek to create high-rank coherent states in the motion of a single trapped ion. The only available measurement is a projective measurement on the *electronic* state of the ion, which is an operator of the form $|e\rangle\langle e| \otimes \hat{\mathbb{I}}_{\text{motion}} = \sum_n |g, n\rangle\langle g, n|$. This cannot be described as a probabilistic projection onto one of a set of states. Instead we must turn to a more general formalism of quantum measurement to test the validity and robustness of the R_3 certifier in this system.

4.3 General measurements

The most general case of operators we will consider are the positive operator-value measures (POVMs) introduced in section 2.3. The interference pattern of eq. (4.4) is generalised to

$$p(\phi) = \text{Tr}[\hat{A}\hat{U}_m\hat{U}_f(\phi)\hat{\rho}\hat{U}_f^\dagger(\phi)\hat{U}_m^\dagger]. \quad (4.11)$$

Without loss of generality, we may replace \hat{A} with a unitary transformation $\hat{U}_m\hat{A}\hat{U}_m^\dagger$ and use the cyclic property of the trace to drop the \hat{U}_m terms. Keeping them separate gives more physical intuition of the process, but is an unnecessary complication for robustness analysis. The convexity of the certifier with respect to the input state is easily shown by the same method that was used previously, so we can immediately progress to calculating the threshold values for different ranks of coherence.

4.3.1 Analytic threshold for 3-coherence

As before, we consider a coherence basis of states $\{|n\rangle\}$ that can be made to evolve by $\hat{U}_f(\phi) = \sum_n e^{-in\phi}|n\rangle\langle n|$. The state and the operator are decomposed in terms of these states as

$$\hat{\rho} = \sum_{n,m} \rho_{nm}|n\rangle\langle m| \quad \text{and} \quad \hat{A} = \sum_{n,m} A_{nm}|n\rangle\langle m|, \quad (4.12)$$

where the coefficients are complex. This form allows the interference pattern eq. (4.11) to be rewritten as

$$p(\phi) = \sum_n \rho_{nn}A_{nn} + 2 \sum_{n>m} |\rho_{mn}A_{nm}| \cos((n-m)\phi + \theta_{nm}), \quad (4.13)$$

in terms of some angles $\{\theta_{nm}\}$ that are the complex phases of the $\rho_{mn}A_{nm}$ terms. All oscillating cosine terms average to zero over the course of one period. This makes the lowest moment, M_1 independent of the relative phases. Similarly, powers of the interference pattern can all be rearranged into sums of terms of the form $\alpha \cos(\beta\phi + \theta_{n_1,m_1} \pm \theta_{n_2,m_2} \pm \dots)$, with $\alpha > 0$ and integer β . The only contributing terms have $\beta = 0$, consequently all moments M_n and the certifier R_3 have a maximum when all the $\{\theta_{nm}\}$ are zero, *i.e.* when $\hat{\rho}$ and \hat{A} are real-symmetric matrices. These are not the only cases when the maximum is reached, but we can proceed under this assumption without loss of generality.

We can analytically calculate the maximal value that R_3 can achieve for any 2-coherent state. With the convexity trivially proven for the general measurement, we need only consider pure states. The energy separation of the two populated

coherence-basis states does not contribute to the value of R_3 , and we will label them $|0\rangle$ and $|1\rangle$. Since the maximum is achieved with a positive real-symmetric density operator, we may parametrise the input state as $\sqrt{x}|0\rangle + \sqrt{1-x}|1\rangle$ for $0 \leq x \leq 1$. Using eq. (4.13), the explicit form of the certifier is

$$R_3 = xA_{00} + (1-x)A_{11} + \frac{6x(1-x)A_{01}^2}{xA_{00} + (1-x)A_{11}}. \quad (4.14)$$

In order for the measurement operator to be a valid value in a POVM, the two on-diagonal elements must have a maximal value of one, and the off-diagonal element must satisfy

$$A_{01} \leq \min\{A_{00}A_{11}, (1-A_{00})(1-A_{11})\}. \quad (4.15)$$

The symmetry in this constraint represents the choice between using \hat{A} or $1 - \hat{A}$ as the measurement operator, so we may examine only the branch with $A_{00} + A_{11} \leq 1$. As expected, R_3 is maximised when the coherence between the two basis states is maximised when the inequality in the constraint is tight, giving $A_{01} = A_{00}A_{11}$. Clearly if any of x , A_{00} or A_{11} are zero, the system described is simply incoherent, and R_3 may attain a maximal value of unity.

The true maximum for 2-coherent states can be found using the method of Lagrange multipliers with the constraints $0 < \{x, A_{00}, A_{11}\} < 1$ and $A_{00} + A_{11} \leq 1$. As only one bound can be tight we need only one slack variable λ , and find

$$\mathcal{L} = R_3 - \lambda(A_{00} + A_{11} - 1) \quad \text{for } \lambda \geq 0. \quad (4.16)$$

The derivative with respect to A_{00} is

$$\frac{\partial \mathcal{L}}{\partial A_{00}} = x \frac{x^2 A_{00}^2 + 2x(1-x)A_{00}A_{11} + 7(1-x)^2 A_{11}^2}{[xA_{00} + (1-x)A_{11}]^2} - \lambda, \quad (4.17)$$

which is transformed into the derivative with respect to A_{11} by the transformations $x \rightarrow 1-x$ and $A_{00} \leftrightarrow A_{11}$. The fraction is strictly positive, so stationary points require that λ is as well, in turn forcing the $A_{11} = 1 - A_{00}$ to satisfy the complementary slackness condition. With all of these conditions, the optimal measurement operator in the restricted $\{|0\rangle, |1\rangle\}$ subspace can be written in matrix form as

$$\hat{A} = \begin{pmatrix} A_{00} & \sqrt{A_{00}(1-A_{00})} \\ \sqrt{A_{00}(1-A_{00})} & 1-A_{00} \end{pmatrix}, \quad (4.18)$$

which is precisely the form of a rank-1 projective measurement of the state $\sqrt{A_{00}}|0\rangle +$

$\sqrt{1 - A_{00}}|1\rangle$. The problem has now been reduced to what was already shown by Dive *et al.*³⁴, and the maximal value of R_3 for 2-coherent states remains $\frac{5}{4}$. Any state that has a measured value of R_3 above this value must be at least 3-coherent.

4.3.2 Numeric evaluation of thresholds

Expanding this direct analysis beyond 2-coherence proves tricky. Instead, we use numerical techniques to maximise the certifier over each space of k -coherent states and general measurement operators, in order to empirically find the upper bounds. At a high level, we simply wish to take any general-purpose maximisation algorithm, such as the quasi-Newton Broyden–Fletcher–Goldfarb–Shanno (BFGS) method⁹⁶, and have it adjust a random POVM value and k -coherent density operator to maximise R_3 . In practice, this means finding a parametrisation of ℓ elements that takes a vector in \mathbb{R}^ℓ smoothly to the search space. Simply taking the individual matrix elements is not suitable, as the constraints on the values become highly non-linear and unsuitable for numerical optimisation. A better way is to craft a parametrisation that is a surjection of \mathbb{R}^ℓ onto the search space; it is permissible—though somewhat undesired—for multiple vectors to correspond to the same pair of POVM value and density operator, but it is required that *all* such pairs have at least one associated parameter vector. If this is achieved, one can use an unconstrained optimisation routine, which are typically orders of magnitude faster and more complete than those for problems with non-linear constraints.

Any POVM value \hat{A} can be written as a sum of simple projective measurements as $\hat{A} = \sum_j a_j |\psi_j\rangle\langle\psi_j|$, for some scalar constants $0 \leq \{a_j\} \leq 1$ and a set of orthonormal states $\{|\psi_j\rangle\}$. For convenience, we will optimise separately over different numbers of non-zero a_j . We approach the parametrisation problem top-down. Each non-zero a_j requires a single parameter, and any standard mapping of $\mathbb{R} \rightarrow [0, 1]$ is suitable for the transformation: logistic transforms, arctangent transforms, and so on. The $\{|\psi_j\rangle\}$ can be chosen by parametrising an orthonormal basis of the complete space, and selecting the desired number from the basis. This can be done by first taking an arbitrary basis of the full Hilbert space, and parametrising a single pure state out of it. The input basis is then limited to cover only the subspace orthogonal to the chosen state, reducing its dimensionality by one. These steps are then repeated, each time removing a dimension from the available subspace for parametrisation by considering only the space orthogonal to all previously selected states. Once enough states have been found, the parametrisation is complete. The reduction of the Hilbert space to the orthogonal subspace can be done via the Gram–Schmidt process, which is suitably deterministic, stable and smooth. A valid

pure state in an n -dimensional Hilbert space can be generated by taking $n - 1$ amplitudes c_q and phases θ_q , and returning the normalised dot product of the vector $(1, c_1 e^{i\theta_1}, c_2 e^{i\theta_2}, \dots)$ with the basis. This shows that there is some duplication in the parametrisation, but in practice it does not pose an issue.

The convexity properties of R_3 should make it unnecessary to draw arbitrary density matrices for the input states, in favour of using pure states, but for completeness' sake we can define a parametrisation. All density matrices that are at most k -coherent can be written as a sum $\sum_j p_j \hat{\rho}_j$, where the $\{p_j\}$ are probabilities and the individual $\hat{\rho}_j$ are arbitrary density matrices each in their own k -dimensional subspace spanned by distinct choices of k basis states from the coherence basis. As with many discrete components in smooth optimisations, we handle the choice of different subspaces by simply repeating the optimisations many times for each possible set of choices. Density operators are positive semi-definite, and thus have a Cholesky decomposition $\hat{\rho} = \hat{L}\hat{L}^\dagger$ for a lower-triangular matrix \hat{L} . We can therefore parametrise an n -dimensional density matrix by drawing $n(n + 1)/2$ parameters to be the magnitudes of the triangular matrix elements and $n(n - 1)/2$ parameters to be the phases of the off-diagonal elements. Matrices parametrised in this way will not give unity-trace density operators, and so if \hat{L}' is the parametrised triangular matrix, the output density operator is $\hat{\rho} = \hat{L}'\hat{L}'^\dagger / \text{Tr}(\hat{L}'\hat{L}'^\dagger)$.

A parametrisation drawn in this manner is clearly biased. Unlike in random sampling where it is strongly preferable to draw from the Haar measure to avoid sampling artefacts, this is not a particular problem for optimisation. The only consideration is to ensure that the optimisation landscape does not become too flat for convergence. This is most likely to be an issue if the optimal value is achieved when certain input parameters must become close to infinite to represent the desired value in the output space, and the target cost varies slowly with respect to large changes in the inputs. In this case, this issue does not frequently arise since the landscape is well featured and the convergence criteria can reliably be reached.

To locate the threshold values with general measurement operators, several thousand optimisations were run in parallel⁵⁰ in Hilbert spaces of varying dimensions, taking varying ranks of the measurement operator \hat{A} and density operators with various numbers of k -coherent components on different subspaces. In all cases, the quasi-Newton method would reduce the total density matrix to a single pure state, setting all but one component probabilities p_j to zero, and create a measurement operator \hat{A} that was precisely a projective measurement onto the input state. This is exactly consistent with the expected results, and the threshold values of R_3 were in total agreement with the prior work³⁴. Further, fixing \hat{A} to be a higher-rank

projector by requiring multiple $a_j = 1$ always resulted in a maximal value of R_3 that was lower than before. Specifically, each additional rank of projector reduced the maximal value achievable by R_3 to the next highest threshold value, for example a rank-2 projector in a 4-dimensional Hilbert space could never measure more than 3-coherence, no matter if the input state was greater. This is intuitive; adding extra orthogonal components reduces the distinguishability of states, which is a key component of the certification.

We have now shown that R_3 is a robust coherence certifier, even for the most general class of quantum measurements. No matter how imperfectly the mapping $\hat{\mathcal{U}}_m$ is implemented, one can never measure a value of R_3 above certain thresholds if the input state does not exhibit sufficient-rank coherence. The only assumption is that any error in the mapping sequence is independent of the free-evolution phase ϕ being applied. In practice, this is easily satisfied in all systems of interest.

4.4 State-creation sequences

The structure of our coherence-creation experiment is a generalisation of the standard two-state Ramsey experiment. In the usual form, a single ion is ideally prepared in the $|g, 0\rangle$ state, although in practice the motion typically has some small thermal component with mean phonon occupation $\bar{n} \ll 1$. A superposition state $|g, 0\rangle + |g, 1\rangle^*$ is prepared by first applying a $\pi/2$ pulse on the carrier—creating the state $|g, 0\rangle + |e, 0\rangle$ —followed by a π pulse on the first red sideband. This is allowed to evolve its phase for some time, then subject to the inverse of the state-creation sequence and measurement. If all operations are implemented perfectly, this shows a sinusoidal response of the ground-state population to the phase evolution. These oscillations are maximum amplitude, because the inverse of the creation sequence maps the target state back to $|g, 0\rangle$ and the only orthogonal state that becomes populated during the evolution, $|g, 0\rangle - |g, 1\rangle$, to $|e, 0\rangle$. There are two major problems preventing the obvious generalisation to *create any superposition, evolve it, invert the creation, and measure*: arbitrary state creation is non-trivial, and inversion of the creation sequence will prove to produce an unsuitable interference pattern.

We will first deal with the creation of arbitrary motional superpositions. The two-level Ramsey scheme works by exploiting the non-interaction of the first red sideband with the $|g, 0\rangle$ state. Directly extending this method to create superpositions with higher motional states would require access to the second-order sideband

*Throughout this chapter we will drop the normalisation factors from state descriptions for legibility. All states considered are actually of unit norm; there is no use of unnormalised states.




FIGURE 4.1—The algorithm used to produce arbitrary motional superpositions⁴¹, illustrated creating $|g, 0\rangle + |g, 1\rangle + |g, 2\rangle$. Grey circles represent state occupation with size proportional to population, and the green, red and blue arrows respectively represent the carrier, red- and blue-sideband transitions. (a) Start from the target state. (b) Apply the red-sideband pulse that moves all population of the highest-occupied motional state $|g, 2\rangle$ into $|e, 1\rangle$, keeping track of the effects on other states. (c) Apply a carrier pulse to join the electronic states of the elements with the now-largest phonon state into either $|g, 1\rangle$ or $|e, 1\rangle$. (d) Depending on the previous pulse, apply either the red or blue sideband to reduce the highest motional state in the system down by a further phonon. (e) Finally, combine the population into the true initial state $|g, 0\rangle$. The desired creation sequence is the adjoint of the operation just derived.

transitions and beyond. As described in section 3.3, the ion–laser coupling in typical trap configurations is insufficient to drive these interactions in a reasonable time frame without causing significant off-resonant effects on other transitions. We must use a method that is limited to first-order sidebands, here a minor extension to previous work^{7,41} that uses both red and blue sidebands rather than just the red.

First consider the operation in reverse, that is with the system starting in the target state of $c_0|g, 0\rangle + c_1|g, 1\rangle + \dots + c_n|g, n\rangle$. The aim is to ratchet down the highest occupied motional state until only $|0\rangle$ remains, then combine all population into $|g, 0\rangle$ to produce the initial state. The adjoint of the whole operation will then be the desired creation sequence. This is illustrated in fig. 4.1.

Explicitly, first apply a red-sideband pulse to move all of the population in $|g, n\rangle$ into the state $|e, n-1\rangle$. This has an effect on every element in the superposition as well—except for $|g, 0\rangle$ —that must be tracked. The different motion levels oscillate with their respective coupled state at generally incommensurate frequencies—approximately $\eta\Omega\sqrt{n}$ for the red- and blue-sideband transitions. The total state after the first pulse is some new $c_{g,0}|g, 0\rangle + c_{e,0}|e, 0\rangle + \dots + c_{g,n-1}|g, n-1\rangle + c_{e,n-1}|e, n-1\rangle$ for complex $\{c\}$, where now various excited electronic states are involved. The pulse angle and phase to achieve this with a red-sideband pulse exactly on resonance can be derived from eq. (3.21) as

$$\Omega_{n,n-1}t = \pm 2 \arctan \frac{|c_{g,n}|}{|c_{e,n-1}|} \quad \text{and} \quad \phi = \arg \frac{c_{g,n}}{c_{e,n-1}} + \pi, \quad (4.19)$$

and equivalent operations using the carrier and blue-sideband transitions have similar forms, albeit with modified motional levels. Next combine the populations of the $|g, n-1\rangle$ and $|e, n-1\rangle$ states into one of these two by using the carrier. Depending on whether the ground or excited electronic state is chosen, apply either the red or blue sideband to reduce the motional level again. Repeat these steps until the state reaches some $c'_{g,0}|g, 0\rangle + c'_{e,0}|e, 0\rangle$, then apply a final carrier to reach the initial state $|g, 0\rangle$, and take the adjoint of the whole sequence to find the desired forwards mapping. Creating a motional superposition with a highest occupied phonon number of n requires at most $2n$ pulses.

This algorithm permits a family of solutions. There is a branch point each time the equal-motion population must be consolidated, since it can be pushed into $|g\rangle$ or $|e\rangle$. The arctangent in eq. (4.19) is also a decision point; since the oscillation frequencies are incommensurate between each pair of coupled states, cycling the upper-most population before consolidating it affects the populations in lower motional states, which may allow for shorter pulses elsewhere in the algorithm. In motion-changing transitions, more excited motional states generally oscillate faster than lower ones, according to eqs. (3.23) and (3.24). For some cases, especially with large superpositions, it may be worth spending longer on a fast pulse in order to save time on slower pulses. The dominant experimental concerns in choosing the particular solution are minimisation of total pulse time and, to a slightly lesser extent, limiting the number of different transitions addressed. Most decoherence channel magnitudes are directly related to the total time: motional heating, frequency drift, voltage instabilities, and so forth. Using fewer transitions reduces experimental complexity by minimising the amount of calibration required. In theory it should be possible to cleanly address all transitions once the lasers are calibrated to the qubit frequency, and the modulation of the laser is calibrated to the trap frequency. In practice, though, some components may show non-linear responses to different target frequencies, and the highest-fidelity operations can typically only be achieved after each transition separation frequency and coupling strength are measured individually.

Despite having an infinite number of possible solutions, one can still find the absolute shortest-time solution with a breadth-first search. First, run the algorithm once making arbitrary decisions about which qubit state to combine populations in, and always taking the shortest pulse angle available as a solution to eq. (4.19). The total time taken by this pulse sequence is an upper bound on the time required. Now, restart the algorithm treating the solution space as a lazily constructed tree (in the computer-science, data-structure sense), where the nodes are decision points

and the branches represent particular choices. Traverse the tree breadth first, attaching the system state (the current statevector and cumulative pulse time taken) to each node as it is encountered. If a node requires a pulse time greater than the upper bound that was previously found, it is rejected and no further solutions to eq. (4.19) need to be considered at this layer. This strategy of pruning the tree as it is constructed ensures that the complete, infinite tree never need be built, and the algorithm will eventually terminate having reached a single leaf node that has the minimum possible time.

Our experiment originally only calibrated the carrier and red-sideband transitions, to reduce the complexity of setting up the experiment. It was found that for the particular motional superpositions chosen, the absolute variation of the optimal superposition-creation sequence from the initial bound—chosen to use only the carrier and red-sideband transitions—was on the order of 1% of the total length. Since this was a very minor improvement, the preference to use only two different transitions generally won out. The explicit forms of the sequences used in the experimental realisation are given later in tables 4.1 to 4.3, along with details of the measurement-mapping sequences that are derived in the next section.

4.5 Measurement-mapping sequences

While section 4.3 showed that imperfect measurements cannot produce false positives, one must still choose a suitable mapping sequence $\hat{\mathcal{U}}_m$ for a given input state to gain the greatest chance of registering a true positive result. To illustrate, no matter how faithfully the mapping and measurement $\hat{A} = |0\rangle\langle 0|$ is implemented, it is entirely incoherent and its resulting interference pattern will always be constant. From the numerical work of the previous section, if the input state $|\psi\rangle$ is an equal superposition of coherence-basis elements, the aim is always to implement a measurement operator \hat{A} that is a rank-1 projector $|\psi\rangle\langle\psi|$. The standard two-state Ramsey experiment achieves this ideal measurement mapping for its input state $|g, 0\rangle + |g, 1\rangle$ by a simple inversion of the preparation sequence. In fact, if the available measurement in a system is a projection onto *exactly* the initial state (before $\hat{\rho}$ is created), inverting the creation sequence is generally sufficient to achieve a high-fidelity mapping for that state.

This simplicity does not extend to the ion-trap measurement $|e\rangle\langle e| \otimes \hat{\mathbb{I}}_{\text{motion}}$ for general states. Measurement of the electronic state in the ion trap does not just measure population in the perfect initial state, but also other motional states. The superposition-creation algorithm described in section 4.4 is specifically constructed

to map $|g, 0\rangle$ to the target state, but its effect on states orthogonal to the target state is not considered at all. In all probability, the orthogonal states that become populated during the phase evolution will be mapped back to have some population in the ground electronic state, and some in the excited. The final measurement of the electronic state is therefore not able to completely distinguish the different cases, and the visibility is limited. Visibility not a direct component of the certifier R_3 , but a lack of it does qualitatively indicate that the maximal achievable R_3 value for this pattern, even implemented perfectly, will be lower than it could be.

While the ion-trap measurement will always fail to distinguish different motional levels, it can be turned into a proxy for a perfect projective measurement by choosing $\hat{\mathcal{U}}_m$ to map the target state into $|e\rangle$, and other states to $|g\rangle$. Let us take an experiment to create $|g, 0\rangle + |g, 1\rangle + |g, 2\rangle$ and verify 3-coherence as an explicit example. We need only consider the space of states that can become populated as a result of the phase evolution. This leads to a small set of conditions for $\hat{\mathcal{U}}_m$:

$$\begin{aligned}\hat{\mathcal{U}}_m(|g, 0\rangle + |g, 1\rangle + |g, 2\rangle) &\propto |e, \lambda_1\rangle, \\ \hat{\mathcal{U}}_m(|g, 0\rangle - 2|g, 1\rangle + |g, 2\rangle) &\propto |g, \lambda_2\rangle, \text{ and} \\ \hat{\mathcal{U}}_m(|g, 0\rangle - |g, 2\rangle) &\propto |g, \lambda_3\rangle.\end{aligned}\tag{4.20}$$

States with more than two phonons do not need to be considered in this mapping, since they do not become populated as part of the state creation or evolution. The choice of the particular other orthogonal states is unimportant, provided the states chosen span the same space as $\{|g, 0\rangle, |g, 1\rangle, |g, 2\rangle\}$; by linearity of $\hat{\mathcal{U}}_m$, any state in this space orthogonal to the target is a linear combination of the two other states chosen, and so will also be purely in the electronic ground state.

Any choice of the motional states $\{|\lambda_j\rangle\}$ —no matter their amount of coherence—is suitable, and equally efficient if the mapping is implemented ideally; the motion is completely traced out and plays no further part. We still must attempt to minimise the total time of this pulse sequence, though, to avoid the effects of the same decoherence processes described in section 4.4. Frequency and phases drifts will cause the manipulations to be imperfectly applied, with longer experimental times leading to a higher likelihood of failing to create or verify the coherence. Motional heating from the trap electrodes or from sympathetic heating from other motional modes cause the phonon numbers to spread out during the operation, naturally destroying the coherence. As with entanglement witnesses, the more these processes occur, and the greater the distance of the actual created state from the state the mapping sequence was designed for, the less concrete information is likely to be obtained

from the method.

The coupled conditions of eq. (4.20) do not appear to permit an analytic solution. Suitable sequences can still be found by numerical methods. The scheme here is far simpler than the complicated parametrisation used in section 4.3 to find the threshold values of R_3 . We start with a list of possible sequences of transitions, such as:

- red, carrier, red, carrier, red;
- carrier, red, blue, carrier, blue;
- blue, red, carrier, blue, red.

Each possibility is optimised for separately. The choices can be of different lengths; empirically, it seems that for a good mapping one requires at least one more pulse than the corresponding state-creation sequence, and these typically alternate between a sideband and the carrier.

For each pulse sequence, each contained pulse is parametrised by two values: the length of time it is applied for, and its relative phase offset. An appropriate loss function is the average population that will end up in the opposite electronic state to where it should. Explicitly, for a target motional superposition state $|\psi\rangle$ and its orthogonal states $\{|\chi_j\rangle\}$, the minimisation problem is

$$\min_{t, \phi} \text{Tr} \left[(|g\rangle\langle g| \otimes \hat{\mathbb{I}}_{\text{motion}}) \hat{\mathcal{U}}_{\text{m}} |\psi\rangle\langle\psi| \hat{\mathcal{U}}_{\text{m}}^\dagger + \sum_j (|e\rangle\langle e| \otimes \hat{\mathbb{I}}_{\text{motion}}) \hat{\mathcal{U}}_{\text{m}} |\chi_j\rangle\langle\chi_j| \hat{\mathcal{U}}_{\text{m}}^\dagger \right], \quad (4.21)$$

where

$$\hat{\mathcal{U}}_{\text{m}}(t, \phi) = \hat{\mathcal{U}}_{\text{m},n}(t_n, \phi_n) \cdots \hat{\mathcal{U}}_{\text{m},2}(t_2, \phi_2) \hat{\mathcal{U}}_{\text{m},1}(t_1, \phi_1), \quad (4.22)$$

and the individual $\{\hat{\mathcal{U}}_{\text{m},j}\}$ are the evolution operators of each transition.

For any realistic experimental realisation, it is not necessary that the loss function exactly reaches zero. I somewhat arbitrarily used a cut-off of requiring the outside-target-state probability to be less than 10^{-10} , and counted any optimisation result within this bound as a success. Each considered sequence of transitions was minimised with a quasi-Newton method equipped with an analytic calculation of the Jacobian of the loss function.* Separate runs were started from random initial parameter vectors, repeating the process for about three hours per sequence. For each target state considered in this work, there were many possible sequences with the same number of transitions used, and within each, several possible values of the parameter vector that qualified as a success. The *best* sequence depends, as

*This is easily, if tediously, calculated from eqs. (4.21) and (4.22) using derivatives of eq. (3.21).

before, on the particulars of the experiment, but in general it is sensible to choose the sequences with the least total duration, and potentially only use the carrier and one other sideband to reduce calibration requirements.

The exact sequences used in our experimental realisation of this certifier are given in tables 4.1 to 4.3, along with the parameters of the state creation. These are also available in a more machine-readable format⁷⁰. In the tables, the *pulse length* is the time apply the pulse, scaled such that a value of 1 is the time taken to completely exchange the populations of the lowest-coupled motional levels in the transition: $|g, 0\rangle \leftrightarrow |e, 0\rangle$ for the carrier, $|e, 0\rangle \leftrightarrow |g, 1\rangle$ for the red sideband and $|g, 0\rangle \leftrightarrow |e, 1\rangle$ for the blue sideband. Note that the pulse lengths on the carrier have significantly less effect on the total duration than those on a sideband, since the power of sideband transitions is suppressed by a factor of the Lamb–Dicke parameter, typically held at around $1/10$. The *phase offset* has the same meaning as it does in eq. (3.21), namely that the driving field phase be offset by this amount relative to where it would have been had it oscillated freely on resonance since the start of the experiment. The phase offsets are not cumulative; each is relative to the start of the experiment.

We were not able to determine a rigorous structure to the pulse sequences found, but there are some features of note. The optimiser strongly preferred sequences that alternate between applying a sideband and applying the carrier. Its first pulse was very frequently the inverse of the last pulse in the state-preparation sequence, which reduces the maximum number of phonons in the system by one. This pulse also populates the majority of the states with lower motional occupation and both a ground-state and excited-state ion. Qualitatively, this allows subsequent pulses to effect more interactions because different states. When the optimiser uses sideband pulses after the initial drive, it often seems to choose a length that would cause the largest motional state to do a complete population cycle with the greater unpopulated state it is coupled to. For example, if the highest motional level at an intermediate point was $|e, 1\rangle$, as in tables 4.1 and 4.2, the applied red sideband would tend to have a duration of $\sqrt{2} \approx 1.41$, which transfers the population of $|e, 1\rangle$ into $|g, 2\rangle$ and then back again, leaving the total phonon number unchanged. The other coupled states do not share the same period, so this appears to be a method by which the optimiser modifies the relative populations in each level, without increasing the maximal excitation.

Beyond these observations, the structure of the optimised pulse sequences appears relatively opaque. The actual motional states that the optimiser maps the different elements to do not have any clear significance. As an example, when

Transition	State creation				Measurement mapping				
	carrier	red	carrier	red	red	carrier	red	carrier	red
Pulse length	0.60	0.80	0.74	0.71	0.71	0.44	1.41	0.54	1.41
Phase offset $/\pi$	0	-0.50	0	-0.50	0	-0.66	-0.83	-0.87	-0.41

TABLE 4.1—Pulse sequences for creation and measurement mapping of target state $(|g, 1\rangle + |g, 2\rangle)/\sqrt{2}$. This is the first non-trivial two-element motional state; superpositions of $|0\rangle$ and either $|1\rangle$ or $|2\rangle$ can be created with only two pulses, and an optimal mapping sequence is just the inverse of the creation. The state here requires a stricter mapping sequence to produce a full-visibility interference pattern. The meanings of the rows are explained in more detail in table 4.2.

Transition	State creation				Measurement mapping				
	carrier	red	carrier	red	red	carrier	red	carrier	red
Pulse length	0.50	0.70	0.73	0.71	0.71	0.48	1.42	1.58	0.71
Phase offset $/\pi$	0	-0.50	1.00	0.50	0	-0.28	-0.25	-0.86	-0.47

TABLE 4.2—Pulse sequence for creation and measurement mapping of target state $(|g, 0\rangle + |g, 1\rangle + |g, 2\rangle)/\sqrt{3}$. The pulse length is the duration of the pulse, scaled such that a value of 1 would completely exchange the populations of the coupled pair of states with the lowest motional occupation. The given phase is applied as an offset relative to where the driving field would have been, had it been oscillating freely since the beginning of the experiment.

	State creation						Measurement mapping			
Transition	carrier	red	carrier	red	carrier	red				
Pulse length	0.51	0.55	0.96	0.57	0.84	0.58				
Phase offset $/\pi$	0	-0.50	-1.00	0.50	0	-0.50				
Transition	red	carrier	blue	carrier	red	carrier	red	carrier	red	
Pulse length	2.89	1.47	1.15	3.02	2.31	4.69	2.31	0.72	0.58	
Phase offset $/\pi$	0	-0.16	-0.41	-0.53	0.45	0.79	-0.32	-0.13	0.76	

TABLE 4.3—Pulse sequences for creation and measurement mapping of target state $(|g, 0\rangle + |g, 1\rangle + |g, 2\rangle + |g, 3\rangle)/2$. This four-element superposition state involved optimising over 18 parameters in an optimisation landscape with many loss-function minima. The long durations of some of the pulses in the measurement-mapping component suggest that there may have been solutions with less time requirements to be found, had we had more time to run optimisations on the compute clusters. The meanings of the rows are explained in more detail in table 4.2.

performing the measurement mapping for the state $|g, 0\rangle + |g, 1\rangle + |g, 2\rangle$ using the sequence given in table 4.2, the different motional superpositions $\{|\lambda_j\rangle\}$ in eq. (4.20) are

$$\begin{aligned} |\lambda_1\rangle &\approx |0\rangle, \\ |\lambda_2\rangle &\approx (0.43 + 0.11i)|0\rangle + 0.75|1\rangle + (0.44 - 0.23i)|2\rangle, \text{ and} \\ |\lambda_3\rangle &\approx 0.72|0\rangle - (0.26 + 0.25i)|1\rangle + (0.05 + 0.59i)|2\rangle, \end{aligned} \quad (4.23)$$

up to global-phase equivalence. It is interesting that the target state is mapped back to the initial motional state, especially because the complete state is *not* the initial state, but is instead $|e, 0\rangle$. For the purposes of certification of coherence, however, we do not need to understand the inner machinations of the optimiser, but simply to use its mappings to implement a perfect projective measurement.

4.6 Statistics of the certifier

One final, major obstacle to the robustness of the R_3 certifier is the impact of imperfect measurement statistics on the interference pattern. So far, the theory work has assumed that the value of R_3 can be determined accurately, and within this assumption it cannot return a false positive for higher-order coherence. This does not survive contact with the real world, however. When measuring an interference pattern, one samples each point by taking a series of *click/no-click* shots and using their count to estimate the probability of an underlying binomial distribution. This provides two mechanisms by which an interference pattern can appear to exhibit coherence that is not present: one in the sampling statistics of individual distributions, the other in the approximation of the continuous interference pattern by a series of discrete points.

As an extreme example, consider the case the state $\hat{\rho} = |0\rangle\langle 0| + |1\rangle\langle 1|$. This is completely incoherent and its interference pattern should be a constant value of $\frac{1}{2}$, but misfortune in the binomial sampling could show false evidence of oscillation. One must account for the standard errors of the measurements when calculating the uncertainty in the value of R_3 , but the ratio of moments of the interference pattern is highly non-linear, meaning the standard physicists' workhorse

$$\alpha_f \approx \sqrt{\sum_j \left| \frac{\partial f}{\partial x_j} \right|^2 \alpha_{x_j}^2}, \quad (4.24)$$

is not valid. Further, the non-linearity introduces a systematic upwards bias in the most natural estimators of R_3 , which must be corrected.

We must first begin with the basics: accurately assessing the uncertainty in individual points of the interference pattern. Each datum is modelled by a binomial distribution with some exactly underlying probability μ_j and a number of shots n taken. The probability of excitation of point j is then a random variate $P_j \sim B(n, \mu_j)/n$ from which our estimates p_j are drawn. The unbiased maximum-likelihood estimator of μ_j is very naturally to define p_j to be the number of successes observed divided by the number of shots. Continuing this, the first statistical moments of the binomial distribution have unbiased estimators (symbols with hats)²¹

$$\begin{aligned} E[(P_j - \mu_j)^1] &\rightarrow p_j - \hat{\mu}_j = 0 && \text{(mean)} \\ E[(P_j - \mu_j)^2] &\rightarrow \hat{\sigma}_j^2 = \frac{p_j(1 - p_j)}{n - 1} && \text{(variance)} \\ E[(P_j - \mu_j)^3] &\rightarrow \hat{\kappa}_j = \frac{p_j(1 - p_j)(1 - 2p_j)}{(n - 1)(n - 2)} && \text{(skewness),} \end{aligned} \quad (4.25)$$

where $E[X]$ is the expectation of a random variable X .

One should not mistake an estimator of the standard deviation of a non-normal statistical distribution for the best estimator of one's confidence in its mean, however. This is commonly done for the binomial distribution, leading to a zero-width interval if zero successes or failures are measured in the finite number of samples, which is in part necessary to avoid a confidence interval that includes invalid values. A more appropriate estimation is the Wilson score interval¹³⁶. This is also based on a normal approximation, but using a slightly different approach, given a number of observed successes k and some desired z score. The standard estimator answers the question *if the distribution has a mean of p_j , what is the expected variance?* while the Wilson score answers *for what range of μ_j would k successes be within the expected interval?* Explicitly, the bounds—without a continuity correction—are the solutions to the quadratic equation

$$\left(1 + \frac{z^2}{n}\right)p_{j,w}^2 - \left(2p_j + \frac{z^2}{n}\right)p_{j,w} + (p_j^2) = 0, \quad (4.26)$$

where the quantities in brackets are known. This interval is not centred on p_j , but is asymmetric; this represents the greater variance that binomial variables with mean close to $\frac{1}{2}$ possess. All results involving binomial variables presented in this thesis use this method for estimating the uncertainty.

We now move to the larger question of estimating R_3 from real experimental data. The integrals in the measurement must be discretised. They vary sufficiently smoothly that a trapezium rule over the J different points in the pattern is appro-

priate. This is

$$\frac{1}{2\pi} \int_0^{2\pi} f(x) dx \approx \sum_{j=0}^{J-1} w_j f\left(\frac{2\pi j}{J-1}\right), \quad \text{where } w_j = \begin{cases} \frac{1}{2(J-1)} & j = 0 \text{ or } j = J-1 \\ \frac{1}{J-1} & \text{all other } j, \end{cases} \quad (4.27)$$

leading to the observed values of R_3 being drawn from a random variable

$$R_3 \sim \frac{\sum_j w_j P_j^3}{\left(\sum_j w_j P_j\right)^2}. \quad (4.28)$$

As alluded to previously, the expectation of this estimator is systematically biased away from the *true* value that would be obtained if the μ_j were known with certainty.

We can analytically evaluate the expectation $E[R_3]$ to determine its bias. The expectation from each point is $E[P_j] = E[\mu_j + (P_j - \mu_j)]$. This unusual form is to permit a Taylor expansion around the binomial means in terms of $E[(P_j - \mu_j)^n]$, which have known unbiased estimators in eq. (4.25). The expansion proceeds, up to terms of third order, as

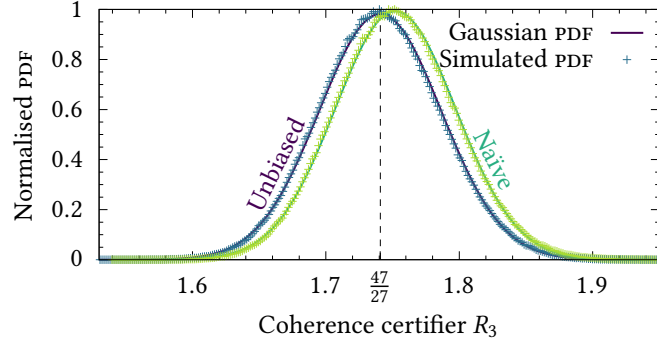
$$\begin{aligned} E[R_3] &= E\left[\left(\sum_j w_j P_j^3\right)\left(\sum_j w_j P_j\right)^{-2}\right] \\ &\approx \frac{\tilde{m}_3}{\tilde{m}_1^2} + \frac{1}{\tilde{m}_1^2} \sum_j \left(3w_j \left(\mu_j - \frac{2}{\tilde{m}_1} \mu_j^2 + \frac{\tilde{m}_3}{\tilde{m}_1^2} \right) E[(P_j - \mu_j)^2] \right. \\ &\quad \left. + w_j \left(1 - \frac{6}{\tilde{m}_1^2} w_j \mu_j + \frac{9}{\tilde{m}_1^2} w_j^2 \mu_j^2 - 4 \frac{\tilde{m}_3}{\tilde{m}_1^3} w_j^2 \right) E[(P_j - \mu_j)^3] \right), \end{aligned} \quad (4.29)$$

bias term

where $\tilde{m}_1 = \sum_j w_j \mu_j$ and $\tilde{m}_3 = \sum_j w_j \mu_j^3$. The actual target of the estimator is the quantity $\tilde{m}_3/\tilde{m}_1^2$. We align our estimator of R_3 to the centre of the distribution by subtracting the marked bias term.

This new estimator is now *fair*, in that it is as likely to return a value that is too great as it is to return one that is too small. We still must calculate our estimate of the uncertainty in the value of R_3 , however. In this case, I found from a multitude of Monte-Carlo testing that the low-order *propagation of uncertainty* formula of eq. (4.24) produces perfectly acceptable results. This is somewhat expected; direct measurements of R_3 appear to be well approximated by a normal distribution, and there is no covariance between the separate p_j . We must then evaluate the partial derivatives. Writing the unbiased estimator as $R_{3,\text{est}} = R_{3,\text{direct}} - \sum_j z_{2,j} - \sum_j z_{3,j}$, where the $\{z_{n,j}\}$ are the terms in eq. (4.29) that include $E[(P_j - \mu_j)^n]$, the derivatives

FIGURE 4.2—Probability density functions for the biased and unbiased estimators of R_3 using a 31-point trapezium rule with 100 shots per point, scaled to have a maximal value of one. The PDFs (crosses) are derived from one million Monte-Carlo simulations of measuring a pattern whose true R_3 is $47/27$, with both estimators using the same set of data. Each PDF is overlaid on a Gaussian approximation (lines) using the estimated values of mean and uncertainty. It is clear that the naïve estimator systematically overestimates.



are

$$\begin{aligned} \frac{\partial z_{2,k}}{\partial p_j} = & \frac{p_k(1-p_k)}{n-1} \left(\frac{8w_j w_k^2 p_k^2}{\tilde{m}_1^4} + \frac{9w_j w_k^2 p_j^2}{\tilde{m}_1^4} - \frac{6w_j w_k p_k}{\tilde{m}_1^3} - \frac{12w_j w_k^2 \tilde{m}_3}{\tilde{m}_1^5} \right) \\ & + \delta_{jk} \left[\frac{p_k(1-p_k)}{n-1} \left(\frac{3w_k}{\tilde{m}_1^2} - \frac{12w_k^2 p_k}{\tilde{m}_1^3} \right) + \frac{1-2p_k}{n-1} \frac{3w_k}{\tilde{m}_1^2} \left(p_k - \frac{2w_k p_k^2}{\tilde{m}_1} + \frac{\tilde{m}_3 w_k}{\tilde{m}_1^2} \right) \right] \end{aligned} \quad (4.30)$$

for the second-order correction terms, and

$$\begin{aligned} \frac{\partial z_{3,k}}{\partial p_j} = & \frac{p_k(1-p_k)(1-2p_k)w_j}{(n-1)(n-2)} \left(-\frac{2w_k}{\tilde{m}_1^3} + \frac{20\tilde{m}_3 w_k^3}{\tilde{m}_1^6} + \frac{24w_k^2 p_k}{\tilde{m}_1^4} - \frac{36w_k^3 p_k^2}{\tilde{m}_1^5} - \frac{12w_k^3 p_j^2}{\tilde{m}_1^3} \right) \\ & + \delta_{jk} \frac{w_k}{\tilde{m}_1^2} \left[\left(\frac{18w_k^2 p_k}{\tilde{m}_1^2} - \frac{6w_k}{\tilde{m}_1^2} \right) \frac{p_k(1-p_k)(1-2p_k)}{(n-1)(n-2)} \right. \\ & \quad \left. + \left(1 - \frac{6w_k p_k}{\tilde{m}_1^2} + \frac{9w_k^2 p_k^2}{\tilde{m}_1^2} - \frac{4\tilde{m}_3 w_k^2}{\tilde{m}_1^3} \right) \frac{6p_k^2 - 6p_k + 1}{(n-1)(n-2)} \right] \end{aligned} \quad (4.31)$$

for the third-order corrections.

The validity of these estimators are tested by Monte-Carlo simulation. We simulate the measurement of an interference pattern generated by an idealised realisation of the superposition state $|0\rangle + |1\rangle + |2\rangle$, including a perfect projective measurement onto this exact target state. The resulting value of the R_3 metric should be exactly $47/27$. The estimated value and uncertainty in R_3 are calculated for one million different attempts, for both the original and unbiased forms of the estimators, each using the same set attempts. The integral discretisation was done with a 31-point trapezium rule, and each point of the pattern was sampled one hundred times. From this, one can derive the probability density function (PDF) of the results by binning the measured values into a histogram of suitable resolution.

The results of these simulations are shown in fig. 4.2. It is clear that the naïve

estimator has a quantifiable bias, and our modifications remove this. Each PDF is shown in comparison to a normal distribution, centred on the mean of the measured values of R_3 , showing the validity of the normal approximation to the uncertainty calculation. On very close inspection, there appears to be a very small skewness in the distribution of the estimated values, some of which is to be expected from the truncation of the Taylor expansion when creating the estimator.

With the exception of trifling details such as *a functioning, controllable ion trap*, we now have all the ingredients in place to create arbitrary motional superpositions, evolve them, and verify that we successfully created higher-order coherence.

4.7 Experimental realisation

Over the course of my studies, the experimental ion-trapping group at Imperial built a new linear radio-frequency trap with a single-segment bladed design. The details of this are given in the thesis of Ollie Corfield²⁸—I was not involved in its construction. This section describes the creation of several-element superpositions in the motional state of a single trapped ion, and their subsequent verification of coherence rank using the methods of the previous sections. I was not part of the lab team implementing the experiment, but I was responsible for interpreting the data; I performed the fitting of the blue-sideband Rabi experiment and of the most likely 3-coherent state to the final results.

The particular ion was $^{40}\text{Ca}^+$, with the two qubit states encoded as an optical qubit with $|g\rangle = 4^2S_{\frac{1}{2}, m_j=-\frac{1}{2}}$ and $|e\rangle = 3^2D_{\frac{5}{2}, m_j=-\frac{1}{2}}$. This was addressed by a single sub-kHz-linewidth diode laser with a wavelength of approximately 729 nm giving a Rabi frequency $\Omega \approx 90 \text{ kHz} \cdot 2\pi$ on the carrier, and readout by the fluorescence measurement shown in fig. 3.2. The initial-state ($|g, 0\rangle$) preparation sequence has a probability of success of 98(2)% through Doppler and then sideband cooling, while the measurement fidelity is reliably above 99%. In terms of the system Hamiltonian from eq. (3.15), the qubit-separation frequency is $\omega_{eg} \approx 411 \text{ THz} \cdot 2\pi$ and the motional frequency is $\omega_z \approx 1.1 \text{ MHz} \cdot 2\pi$, giving a Lamb–Dicke parameter $\eta \approx 0.09$. This is well within the Lamb–Dicke regime, and as such the higher-order sidebands are not reasonably available. With the sideband transitions being suppressed by around an order of magnitude compared to the carrier, there is a sizeable AC Stark effect shifting the frequencies of these transitions. To mitigate this to first order, an additional compensation pulse is applied far off-resonantly, halfway between the carrier and the opposite-colour sideband.

4.7.1 State creation

We attempted to create three different motional superpositions, and verify the coherence rank of each. The three states were $|1\rangle + |2\rangle$ for a 2-coherent state, then $|0\rangle + |1\rangle + |2\rangle$ for a 3-coherent state and $|0\rangle + |1\rangle + |2\rangle + |3\rangle$ for 4-coherence. For 2-coherence, $|1\rangle + |2\rangle$ was chosen as it is the first state that requires the non-trivial state-creation and measurement-mapping sequences to generate it. Note that the relative phases of the superposition elements are important in both parts of the process. We kept all elements the same phase for simplicity in describing them.

After the state creation, one can take an incoherent measure of the populations of each state as a sanity check. All population should be in the electronic ground state of the ion, and this can be tested directly with the standard ion-trap measurement. Due to the nature of the superposition-creation algorithm, having all population in $|g\rangle$ is a very good indication that the algorithm was successful; miscalibrated frequencies, pulse powers or phase drift would all cause appreciable population to remain in $|e\rangle$ at the completion of the algorithm.

We can also make an estimate of the populations in each motional level with a Rabi-type experiment. The superposition state is evolved by applying the blue sideband for a varying amount of time. Each motional level $|g, n\rangle$ will undergo oscillation between the ground and excited electronic states at a rate proportional to $\sqrt{n+1}$ inside the Lamb-Dicke regime. Scanning the length of the pulse builds up a pattern such as fig. 4.3 for the three-element superposition.

One can approximately determine the populations of the different basis states from this pattern. Ideally, we would take a Fourier transform of the data and examine the amplitudes of the different components to find the populations. This is difficult in practice, however. Since the component oscillation frequencies are 1, $\sqrt{2}$, $\sqrt{3}$, and so on, any discrete Fourier transform will not have bins centred on these values, and there is significant leakage between the frequencies unless the scan can last far longer. The coherence time of the motion, while not directly measured, limits this scan from being taken much beyond the 1 ms shown in fig. 4.3.

Instead, we can estimate the populations by a maximum-likelihood method. In essence, one parametrises a model, then finds the values for which the resulting interference pattern would have the greatest probability of returning the observed data. This likelihood is the product

$$\ell(\mathbf{x}) = \prod_j \Pr[B(n, p_j(\mathbf{x})) = n_j], \quad (4.32)$$




FIGURE 4.3—Excited-state population of state $|g, 0\rangle + |g, 1\rangle + |g, 2\rangle$ while being driven by the blue sideband. The data points (purple crosses) are shown with their Wilson binomial $1\text{-}\sigma$ confidence interval. The pattern from a state that best fit the data (blue line) was found by maximum-likelihood estimation, and its 95% confidence region (blue shaded region) by bootstrapping the measured data 14 000 times. The fit accounted for the Rabi frequency, detuning, motional dephasing rate, basis-state populations up to $|3\rangle$ inclusively, and correlations between directly coupled elements. The populations in $|g, 0\rangle$, $|g, 1\rangle$ and $|g, 2\rangle$ were 33(2)%, 30(2)% and 33(2)%, respectively, with 4.7(14)% elsewhere. All appreciable undesired population was in the $|e\rangle$ excited qubit state; the motional state $|3\rangle$ was included in the fit, but found to have a population consistent with zero with a standard error of 9×10^{-3} percentage points.

where n_j is the number of times the measurement returned $|e\rangle$ at a given point, and p_j is the simulated probability at that point for a parameter vector \mathbf{x} . The maximum of this cannot generally be calculated analytically. For numeric evaluation, it is generally advisable to use the log-likelihood $\ln(\ell)$ to avoid floating-point underflow. Doing so has no effect on the location of any maxima as the natural logarithm is a monotonically strictly increasing function over this domain. Any standard optimisation routine can then be used to maximise the quantity.

The parameters to be fit to the data here were: the base Rabi frequency of the sideband; a detuning from the actual transition frequency; the basis-state populations of $|g, n\rangle$ and $|e, n\rangle$, with n up to one phonon larger than the maximum expected; and phase correlations between pairs of states coupled by the blue sideband. In addition, we account for potential motional dephasing by simulating the Lindblad master equation of eq. (2.17) with an additional jump operator $\hat{L} = \sqrt{\gamma}\hat{a}^\dagger\hat{a}$ for some rate γ included in the fit. Physically this additional term is a proxy for the effect of a drifting laser phase. We satisfy the requirement that this operator is bounded by truncating the Hilbert space considered to only the required motional levels. There are no processes being modelled that would cause the population to move outside this subspace.

Performing maximum-likelihood estimation alone does not give any indication of the level of confidence in the fitted values. The simplest method for obtaining an estimate of this is by bootstrapping. In this, the sample of measured data points are

Target state in $ g\rangle$	$\eta\Omega/(2\pi \text{ kHz})$	$\delta/(2\pi \text{ kHz})$	$\gamma/(2\pi \text{ kHz})$
$ 1\rangle + 2\rangle$	8.39(3)	1.3(5)	0.11(2)
$ 0\rangle + 1\rangle + 2\rangle$	6.95(3)	2.6(2)	0.14(3)
$ 0\rangle + 1\rangle + 2\rangle + 3\rangle$	7.52(2)	1.5(3)	0.22(4)

TABLE 4.4—Values and uncertainties in the fit parameters after bootstrapping the measured data for each of the three target states. The modified Rabi frequency $\eta\Omega$ is the coupling strength to the lowest motional state on the blue sideband at zero detuning. Each detuning δ is on the order of $1/500$ of the sideband separation frequency, but closer in size to the Rabi frequency than is desirable. The dephasing strength γ is relatively small compared to all other parameters.

Target state in $ g\rangle$	Estimated population in basis state					
	$ g, 0\rangle$	$ g, 1\rangle$	$ g, 2\rangle$	$ g, 3\rangle$	$ g, 4\rangle$	$ e\rangle$
$ 1\rangle + 2\rangle$	0.052(14)	0.546(14)	0.383(13)	0.000(5)		0.018(9)
$ 0\rangle + 1\rangle + 2\rangle$	0.33(2)	0.31(2)	0.33(2)	0.000(12)		0.036(11)
$ 0\rangle + 1\rangle + 2\rangle + 3\rangle$	0.29(2)	0.25(2)	0.21(2)	0.223(15)	0.000(11)	0.026(11)

TABLE 4.5—State populations, not complex amplitudes, for individual quanta of motion in the ground qubit state, and the total population in the excited qubit state. The values and their uncertainties were found by maximum-likelihood estimation while bootstrapping the measured data. Each state shows some deviation from the ideal, but not excessive. In all cases, the population in motional states larger than those in the superposition was statistically consistent with zero.

repeatedly *resampled* to gain knowledge of how the fit responds to variations in the sampled values. For each new sample, the same maximum-likelihood estimation is performed many times, to give an estimate of the extent of the uncertainty. There are multiple possible methods for performing the resampling, depending on the type of problem. The most straightforward is simply to select a new set of n points from the measured data with replacement, such that some points are sampled more than once and some not at all. Alternatively, one can consider a *parametrised* resampling, where the maximum-likelihood fit is taken to be the underlying pattern, and each new set of data is generated by drawing each point from the binomial distribution implied by the pattern. In the analysis here, I used both of these methods for 180 CPU-hours per method per state, which resulted in approximately 14 000 realisations for each. The 95% confidence intervals for the true pattern from the two methods, as shown by the light region in fig. 4.3 for $|g, 0\rangle + |g, 1\rangle + |g, 2\rangle$, were within statistical noise of each other.

The measured parameters for the three states are shown in tables 4.4 and 4.5. While the Rabi frequencies are consistent with the measured values from calibration, the detunings exceed what was expected. It is plausible that the simulation model was partially underparametrised, such as by failing to account for inefficiencies in

the measurement process or off-resonant excitations from the AC Stark effect, and that this led to better fits with a larger detuning. Maximum-likelihood estimation is, after all, only as good as its model. In this case, the significantly larger uncertainties on the detuning suggest that it had relatively little impact on the model, and there is still good reason to believe the other parameters with much lower relative uncertainties.

4.7.2 Coherence certification

With the motional states created and sanity-checked, all that is left is to attempt to verify the rank of the coherence that has been created. The phase evolution \hat{U}_f of eq. (4.11) is implemented virtually in this system; each transition driven after the period of the phase evolution is offset by an amount $n\phi$, where n is the difference in the number of phonons between the coupled ground and excited qubit states: 0 for the carrier, -1 for the red sideband and 1 for the blue sideband. This is exactly equivalent to evolution under the free Hamiltonian $\hat{\mathcal{H}} \propto \sum_j j|j\rangle\langle j|$ if only the motional states are considered.

For the certification scheme, it is vital that no evolution-phase-dependent error enters the measurement-mapping operation \hat{U}_m . This is a core assumption of the derivations in the threshold values of the certifier. It is possible in theory for a completely incoherent state $|0\rangle$ to be falsely detected as 2-coherent if the mapping sequence somehow projected it onto the state $\cos\phi|0\rangle + \sin\phi|1\rangle$. While this is an extreme example, we still must consider any ways that our system could introduce any phase-dependent error.

The virtual phase advancement on the surface looks plausible, but since the resulting pattern must be periodic, this would be trivially detectable; if the measured interference patterns did not appear to have a period of 2π , it would be clear that some catastrophic failure of the control systems had occurred. This phase advancement is applied in the arbitrary-waveform generator in the same manner as all other phase offsets, and consequently any introduced error is independent of the magnitude of the shift. More detail on how the pulses are actually synthesised is presented in the theses from the experimental group²⁸. In fact, applying the phase advancement virtually has an advantage in that it is a constant-time operation, so there is no chance that run-time-dependent processes will enter as each shot has precisely the same duration.

Beyond this, it is possible that some drifting parameters in the lab could have a time-dependent effect on measurement efficiency or the mapping sequence. With all the cooling and measurement cycles and at 400 shots per point, the final data

collection period for each experiment was around ten minutes. This was in part limited by the length of time the system calibration was valid for; beyond this, and the trap parameters may have drifted far enough that the fidelity of the coherent manipulations was compromised. To avoid any possibility that the environment had an effect that appeared to be phase-dependent, the different shots of the experiment were interleaved. 24 different phase-advancement values were taken, and each was repeated 400 times. Rather than running every shot for one point, and moving on sequentially, a random order of the 24 points was generated, and then this order was cycled through taking one hundred shots per raster. Limitations in the control system made it impractical to randomise the order for each individual shot. Regardless, this has the approximate effect of converting any time-dependent noise from the environment into a white-noise process, which is safely handled by the coherence certifier.

A final possibility is for a systematic mis-set of the driving frequencies. If the qubit frequency is mis-set from the average such that the red- and blue-sideband transitions are addressed at different detunings, any phase advancement on them could in principle introduce a phase-dependent error. This could arise from an imperfect compensation of the AC Stark effect, or by calibrating both sideband frequencies using only a measurement of the qubit frequency and the trap frequency. In this experiment, the two sideband locations were calibrated independently, which alleviates the majority of this concern. The drift of the laser frequency with respect to the sideband frequency over the course of an entire experiment was estimated by pre- and post-calibration to be less than $\delta/\Omega = 0.15$, where the time-dependent components of this will have been converted to incoherent noise by the shot randomisation described previously.

With all this in hand, fig. 4.4 shows the measured interference patterns and ideal models for the two target states $|g, 1\rangle + |g, 2\rangle$ and $|g, 0\rangle + |g, 1\rangle + |g, 2\rangle$. Of these, the two-element superposition had a measured value of $R_3 = 1.090(12)$ after the bias correction, which is above the threshold of 1 officially needed to certify 2-coherence with this metric, though of course the evidence of oscillation alone is sufficient in this case. As expected, the certifier value does not reach 1.25, the maximum achievable with a 2-coherent state, due to imperfections in the experimental environment.

The three-element superposition, on the other hand, did reach above this, to $R_3 = 1.54(2)$. Achieving this value certifies that the state created *must* have contained higher-order coherence; to reiterate, it is impossible for a state without genuine 3-coherence to exhibit a value of R_3 larger than 1.25. For this state, the


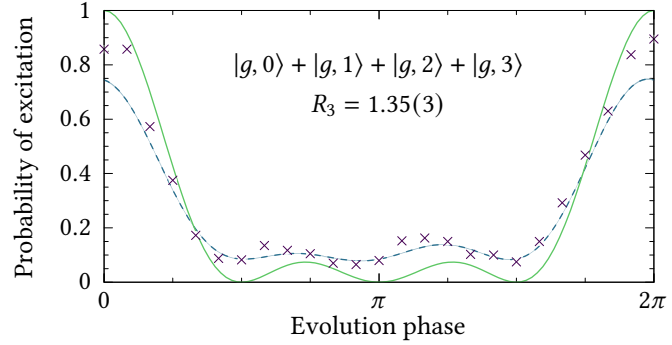


FIGURE 4.4—Measured interference patterns (purple crosses) for the two- and three-element motional superpositions indicated, using the measurement-mapping sequences described in tables 4.1 and 4.2. Each point was repeated 400 times, and the Wilson 1- σ confidence intervals are indicated with error bars though these small enough to be difficult to see. The models (green lines) are the theoretically optimal interference patterns, had there been no experimental error or noise. The measured certifiers were greater than the level needed to certify 2-coherence (1) and 3-coherence (1.25) respectively.

FIGURE 4.5—As in fig. 4.4, but for the four-element superposition indicated with the measurement-mapping sequence described in table 4.3. The value of the certifier here, $R_3 = 1.35(3)$, does not reach the level of $^{179}/_{96} \approx 1.86$ that is necessary to unambiguously certify 4-coherence, but it is at least 3-coherent. The interference pattern of the 3-coherent state that best approximates both this and the state-population data is also shown (dashed blue line), where the thinner, lighter line within is its 1- σ Wilson interval. The unconvincing fit is evidence that the state likely was 4-coherent, but the certifier test was inconclusive.



optimised mapping pulses are absolutely necessary to this certification; without them, the certifier would have a theoretical maximal value of $R_3 \approx 0.92$. This is significantly below the threshold needed to certify, in part because the peak-to-peak visibility of the pattern could not have exceeded 0.68.

Attempting to stretch the trap—which was not specifically designed to have any particular resilience to motional decoherence—to its limits, we also pursued the four-element superposition $|g, 0\rangle + |g, 1\rangle + |g, 2\rangle + |g, 3\rangle$. The resulting interference pattern is shown in fig. 4.5. This state did not cross the threshold of 1.86 necessary to certify 4-coherence, but its value of $R_3 = 1.35(3)$ was still larger than 1.25, unequivocally still certifying it as 3-coherent.

In all three cases, the observed interference patterns were of lower total visibility than they theoretically could have been. Figures 4.4 and 4.5 also show the ideal interference patterns that could have been measured, had the state preparation and measurement mapping been implemented without any experiment error. For the

two-element superposition, this would have been the threshold value 1.25, while for the three-element superposition it would have been $47/27 \approx 1.74$. This is not the absolute maximum achievable for a 3-coherent state, but well above the threshold. The currently known maximum is 1.77, which is exhibited by a superposition with the central element weighted slightly higher than the outer two³⁴. The four-element superposition could potentially have reached $145/64 \approx 2.27$.

It is natural that the highest possible values of R_3 were not observed in this real experiment. The fits of the state populations in table 4.5 show that it is likely the state creation was not exact, and one naturally expects that the measurement-mapping operations, being slightly more complex, would similarly be implemented with some inaccuracies. Frequency drift, off-resonant excitation and thermalisation of the motion all likely played some part in the reduction of fidelity. Of course, some of these processes of sufficient amplitude will destroy higher order coherence. This is where the resilience of the certifier to always fail safe is most important; even in the presence of experimental imperfections, the observed values of the certifier still guarantee that genuine 3-coherence was created in the three- and four-element superposition tests.

For the four-element superposition, which failed to reach a value of R_3 sufficient to classify it as 4-coherent, we can extend the analysis a little further. Using the coherent-state parametrisations and maximum-likelihood techniques introduced in sections 4.3.2 and 4.7.1 we can optimise to find the 3-coherent state that is most likely to have produced the observed data. We do not need to limit ourselves to the final interference pattern in this case; we can also use the blue-sideband-scan data generated during the initial tests of the superposition creation. The results of this best state are plotted as the blue dashed line in fig. 4.5. This does not produce a particularly convincing fit; it is similar towards the centre of the pattern, but has significantly reduced visibility overall. It is perhaps likely, then, that the state created was in fact 4-coherent, but the R_3 certifier returned only an inconclusive result. This is not a failure of the certifier, but more a further example of its fail-safe nature; it will either describe a state as *certainly k-coherent* or not offer an opinion. If one is prepared to relax the burden of proof from *beyond reasonable doubt* to a mere *balance of probability*, the maximum-likelihood estimation offers additional information, without further experimental cost.

4.8 Conclusion

Multilevel coherence is a resource, and despite its complex definitions in terms of density operators, its presence can be certified by one-dimensional interference pattern experiments. This remains true when the coherence basis itself is not accessible to direct measurement, such as for the motional states of a single trapped ion. Even in cases where the certifier test is inconclusive, one can supplement the analysis with maximum-likelihood-estimation techniques to indicate whether it is probable that the observed interference pattern was caused by a state with the particular rank of coherence. The certifier is robust against false-positives, and is massively more practical to implement than complete state tomography, especially in cases with a poor measurement operation.

We have shown that previous interference-pattern methods for coherence certification can be extended to the case of these non-ideal measurement operators, analytically for low ranks of coherence and numerically beyond that. We have performed a detailed analysis of the statistical properties of any measurement of this certifier in an experimental setting, and shown how the bias in the standard estimator of it can be reduced. A scheme for optimal measurement mappings was demonstrated that converted low-visibility patterns into ones that could be used to certify high-rank coherence, and these were demonstrated in a real quantum system in collaboration with the experimental group at Imperial. Two states that unambiguously exhibited properties of 3-coherence were created and verified in the motional state of a single trapped ion. We also provided some additional analysis that indicate higher coherence, when certification is inconclusive.

The intricacies of high-order coherence remain little understood, but this chapter has introduced a new method of probing them. The methods described here are generally applicable far beyond trapped ions. All that is required is a simple measurement and a phase evolution that is very naturally available to any system containing a quantum harmonic oscillator component, such as optomechanical oscillators. This opens up the physical systems that can be used to investigate higher order coherence. Going further, the construction of quantum computers requires incredibly detailed control over large-scale superpositions. A method for generating the necessary phase advancement on an arbitrary number of qubits using only local operations was sketched out, allowing these interference-pattern methods to be used in these situations.

Of course, the other major requirement in quantum computing is the creation of entanglement that is robust to variations in the environment. We now add a

second ion to our theoretical trap, and move to discussing how we can generate a Bell state between them that will not suffer as heavily from environment noise.

Chapter 5

Robust Entangling Gates

COAUTHORSHIP

All the work in this chapter was done by me, with supervision by Florian Mintert. The results here were also published as ref. 72.

We described the Mølmer–Sørensen scheme for generating entanglement in trapped ions in section 3.4. Its development at the time presented a large advance for quantum logic gates in this medium, relaxing the previous requirement²⁵ that the coupled motional modes were cooled perfectly to their ground state¹¹⁵, although it originally came at the penalty of adiabaticity. The more strongly coupled version of the interaction described by eq. (3.28) was developed shortly after¹¹⁶, and is now by far the more common method for applying the gate^{30,40,48,139}. This trades off increased sensitivity to fluctuations in the control frequencies for significantly faster gate operation. Faster gates mean less time in which the quantum systems can decohere, so in practice this exchange is always worthwhile. Still, it is entirely valid to ask *can we reduce the sensitivity to errors?*

This is of course not a new question. The most obvious method of increasing fidelity is to improve the quality of the controlling electronics and drive fields, and to reduce any environmental effects that could cause frequency shifts. Technology on this front is always improving, allowing experimentalists to move from the earliest fidelities of around 80%⁴⁵, to 99% shortly after⁸, to the most modern realisations in excess of 99.9%⁴⁰. Tighter control tolerances are not solely responsible. Recent implementations of these trapped-ion gates all use some additional techniques to eke out more speed, such as addressing multiple motional modes simultaneously^{2,23} and driving Raman transitions closer to resonance to increase the transfer rate⁴.

One can attain even greater fidelity by accepting that noise will always be present, and minimising not just its source but its ability to affect the system. The classical spin-echo technique in nuclear magnetic resonance⁴⁴ is as applicable to modern quantum information processors as it was then, whether in qubit idle periods or during gates, now under the name *dynamical decoupling*¹²⁷. These techniques

are typically applied in microwave-driven gates^{12,49,122}. Beyond this, one can apply optimal-control techniques to shape control fields, to suppress undesired effects at various points in the gate application. There are a variety of different parametrisations for this, from simple smoothing of the pulse windows⁸ to more complex schemes based on piecewise-constant functions¹⁰⁷, or amplitude^{42,110,132} or phase modulation^{80,140}.


Such error-mitigation techniques continue to grow in importance. All proposals to enlarge ion-trap quantum computers inevitably increase the complexity of the systems, whether this is by shuttling ions between modules⁶⁸ or linking traps with photonic interconnects^{84,118}. This leads to greater physical differences and larger quantities of bulk optics over which the control fields must be kept coherent, while attempts to miniaturise traps necessarily lead to greater heating rates and neighbour interactions via crosstalk¹⁹.

This chapter presents an investigation from early in my degree into the applicability of Fourier-series *multi-tone* parametrisations of the Mølmer–Sørensen gate for suppression of mis-sets in the qubit frequencies. The scheme is applicable to any physical encoding of the qubits, whether they are driven by microwaves, a single laser targeting a dipole-forbidden transition, or two lasers in a Raman configuration. It requires no additional fields, only shaping of the existing control; the only necessity over the original implementation of the Mølmer–Sørensen interaction¹⁰³ is an arbitrary waveform generator. We consider only the case of gates that do not increase the required peak power, consistent with realistic experimental considerations. This method, in theory, can improve the infidelity scaling of the gate with respect to the frequency error, given a sufficient number of tones in the control fields. While we had originally hoped to extend a previous experimental collaboration with the group at the University of Sussex¹³², they had moved on to other priorities and we were unable to test it experimentally.

5.1 Model

The Hamiltonian of the trapped-ion system was given in eq. (3.15), under the critical assumption that all qubits had the same frequency. We cannot necessarily assume this in the general case. Instead, let us consider the case of two co-trapped ions, whose frequencies are described by $\{\omega_{eg,i}\}$. In an ideal situation for ions of the same species, the frequencies should be identical for some known ω_{eg} . It is more convenient, then, to think of the two ions as having qubit frequencies slightly detuned from this ideal case, by some small amount $\{\delta_{eg}^{(j)}\}$. We then can

FIGURE 5.1—The effects of different frequency off-sets on the operation of the Mølmer-Sørensen gate. This is more complete version of fig. 3.4, taking into account the new error terms introduced in eq. (5.1). Thick black lines are the ideal energy levels targeted by the driving, while thin lines are the true structure. The qubit errors δ_{avg} and δ_{spl} cause the two-photon process to be off-resonant for some initial states, while a motional error δ_z causes a shift from the ideal virtual level.



more explicitly encode the symmetry of the problem under the exchange of ions by working with a single quantity, the frequency splitting $\delta_{\text{spl}} = (\delta_{eg}^{(1)} - \delta_{eg}^{(2)})/2$ between the two qubits.

This prompts some modifications to the interaction Hamiltonian of eq. (3.19), principally that the sum over the individual qubit operators is no longer the simple $e^{-i\omega_{eg}t} \sum_j \hat{\sigma}_+^{(j)}$, but must take into account the separate frequencies on each ion. Further, as we will be moving to shape the control fields, we replace the simple laser driving with an abstract profile $f(t)$. Doing so allows the possibility that the driving will be defined with respect to miscalibrations of the average qubit frequency and motional frequency as well; we introduce new factors of $\delta_{\text{avg}} = (\delta_{eg}^{(1)} + \delta_{eg}^{(2)})/2$ and δ_z respectively to account for these. This leaves a new interaction Hamiltonian of

$$\hat{\mathcal{H}}_{\text{int}}/\hbar = f(t)e^{i(\omega_{eg}+\delta_{\text{avg}})t} \left(e^{i\delta_{\text{spl}}t} \hat{\sigma}_+^{(1)} + e^{-i\delta_{\text{spl}}t} \hat{\sigma}_+^{(2)} \right) \hat{\mathcal{D}} \left(i\eta e^{i(\omega_m+\delta_m)t} \right) + \text{H.c.}, \quad (5.1)$$

where $\hat{\mathcal{D}}(\alpha) = \exp(\alpha\hat{a}^\dagger - \alpha^*\hat{a})$ is the displacement operator from eq. (2.6). The effects of these on the energy levels for a standard Mølmer-Sørensen gate targeted exactly on ω_{eg} and ω_z is shown in fig. 5.1.

In practice, one must always apply control fields that are close to resonant with particular sidebands to drive meaningful dynamics. We rewrite the most general field $f(t) = \tilde{f}(t)e^{-i\omega_s t}$, re-using the sideband selection frequency $\omega_s = \omega_{eg} + n\omega_z$ discussed in section 3.3.2. The remaining time dynamics in the $\tilde{f}(t)$ are assumed to be slow compared to the sideband separation, and the selection frequency is deliberately defined in terms of the calibration parameters ω_{eg} and ω_z , rather than the true values of the average and the motion frequency, to mimic a real experimental setup. The Mølmer-Sørensen interaction is achieved by applying two component driving fields: $f_r = \tilde{f}^*$ on the red sideband, and $f_b = \tilde{f}$ on the blue. This leads to a complete description of the Mølmer-Sørensen Hamiltonian within the

Lamb–Dicke approximation, in terms of all of these frequency errors, as

$$\hat{\mathcal{H}}_{\text{MS}} = -\eta \tilde{f}(t) e^{i\delta_z t} \hat{a}^\dagger \cdot \begin{pmatrix} \cos((\delta_{\text{avg}} + \delta_{\text{spl}})t) \hat{\sigma}_y^{(1)} \\ + \sin((\delta_{\text{avg}} + \delta_{\text{spl}})t) \hat{\sigma}_x^{(1)} \\ + \cos((\delta_{\text{avg}} - \delta_{\text{spl}})t) \hat{\sigma}_y^{(2)} \\ + \sin((\delta_{\text{avg}} - \delta_{\text{spl}})t) \hat{\sigma}_x^{(2)} \end{pmatrix} + \text{H.c.} \quad (5.2)$$

In the absence of all errors, this degrades to an equivalent form to eq. (3.26) with generalised driving.

Each error is largely caused by a separate miscalibration or experimental imperfection. Qubit-splitting errors can arise when the chosen encoding is magnetic-field sensitive, and there is a field gradient along the axis of the trap. A miscalibration of the average frequency of the qubits can arise from a drift in the laser caused by inadequate locking, or from a global magnetic field drift across all qubits. The motional frequency can often be affected by fluctuations on trap endcap electrode voltages, but it can also be difficult to calibrate this precisely, due to the AC Stark effect from the nearby stronger carrier transition when probing sidebands.

The errors have different effects on the gate operation. An error in the motional frequency affects the red and blue sidebands equally but in opposite directions, meaning that the two-photon red–blue process of the gate is on-resonant overall for all states. This does, however, mean that the desired detuning from the sidebands, the ϵ in eq. (3.28), is not what is expected, and consequently both the applied Rabi frequency and the gate time will be incorrect, leading to residual qubit–motion entanglement at the completion of the pulses, and an incorrect amount of spin-dependent phase advancement. Both of these will strongly affect the quantum information stored in the qubits; in quantum computing applications, the coherence time of the motion is significantly shorter than the coherence time of the qubits.

The two qubit frequencies play a separate role. Qualitatively, from fig. 5.1, it is clear that a nonzero error in the qubit average will cause the two-photon process from an initial state of $|gg\rangle$ to be off-resonant; the energy separation between $|gg\rangle$ and $|ee\rangle$ is different the sum of the two photon energies by a factor of $2\delta_{\text{avg}}$. From this same starting state, the splitting frequency will cause an additional decoherence of the qubits by making them distinguishable during the operation, removing some of the path interference that cancels out the motional dependence in the ideal gate, but will not alone cause the complete process to be off-resonant.

It is important to note, however, that the $|gg\rangle \leftrightarrow |ee\rangle$ process is only part of the story. For complete gate operation, one must also consider the coupling of

$|ge\rangle \leftrightarrow |eg\rangle$. In this, the driven processes do not use a photon from each sideband, but two photons from the same sideband, and thus of the same frequency. This means that the splitting error now causes off-resonant effects, while the average error shifts the virtual levels targeted by the red–red and blue–blue processes by different amounts. Formally, for a physical gate $\hat{\mathcal{U}}$ that is attempting to implement some target dynamics $\hat{\mathcal{U}}_{\text{tg}}$, we are concerned with the *average gate fidelity*

$$F = \frac{1}{J} \sum_j \text{Tr}(\hat{\mathcal{U}}|j\rangle\langle j|\hat{\mathcal{U}}^\dagger\hat{\mathcal{U}}_{\text{tg}}|j\rangle\langle j|\hat{\mathcal{U}}_{\text{tg}}^\dagger) = \frac{1}{J} \sum_j \left| \langle j|\hat{\mathcal{U}}_{\text{tg}}^\dagger\hat{\mathcal{U}}|j\rangle \right|^2, \quad (5.3)$$

where the J different $\{|j\rangle\}$ form a complete basis. In practice, we would always precalculate the $\hat{\mathcal{U}}_{\text{tg}}|j\rangle$ terms only once, to avoid an extra matrix–vector product in an inner loop.

It is clear from eq. (5.2) that no amount of shaping f_b will actually remove terms from the Hamiltonian. Instead, we seek a pulse design that reduces errors in the *effective* Hamiltonian that is applied at the gate time. Unfortunately, with Pauli operators with time-dependent amplitudes now featured, the Magnus-expansion approach used to analytically calculate the dynamics in section 3.4 no longer terminates. The aperiodicity of the system also prevents the frequently used Floquet approach from being valid, and we need to turn to numerics to reasonably simulate the dynamics.

We shall consider a multi-tone parametrisation, which is conceptually similar to several Mølmer–Sørensen gates being applied simultaneously, with each detuned by an integer multiple of the base case. For a driving field with n tones, we write the blue-sideband driving profile as

$$\tilde{f}(t) = \frac{\epsilon_n}{4} \sum_{k=1}^n c_{n,k} e^{-ik\epsilon_n t}, \quad (5.4)$$

for nondimensional complex numbers $\{c_{n,k}\}$, and a known base detuning ϵ_n . The scaling factor of $2\epsilon_n$ is chosen such that $c_{1,1} = 1$ reproduces the original, single-tone Mølmer–Sørensen gate. The value of the detuning determines the gate time, and it is generally limited by the available driving-field power; a larger value requires more power in order to reduce the gate time as $\tau_n = 2\pi/\epsilon_n$.

5.2 Optimisation

One very natural figure of merit for a gate is its average fidelity, as defined in eq. (5.3). As a minor implementation detail, numerical optimisers are traditionally

implemented as minimisers rather than maximisers—one person’s f is another’s $-f$, making a distinction redundant—which motivates targeting the *infidelity* $I = 1 - F$ instead. This is also a physically sound choice; we generally expect infidelity to never reach zero, but its order of magnitude plays a key role in whether quantum error correction is possible⁶⁴.

However, in the absence of errors, and assuming the Lamb–Dicke regime holds perfectly—more on that in the next chapter—the Mølmer–Sørensen gate has zero infidelity. This is no longer the case once errors are present, but we cannot optimise for a specific value of the error; if we knew it, we would mostly be able to recalibrate the system to remove it. In reality, we will have some modelled probability distribution of the errors $\omega(\boldsymbol{\delta})$ and we care about reducing the expectation of the infidelity over all possible errors:

$$\mathbb{E}[I] = 1 - \frac{1}{J} \sum_j \int d\boldsymbol{\delta} \omega(\boldsymbol{\delta}) \left| \langle j | \hat{U}_{\text{tg}}^\dagger \hat{U}(\boldsymbol{\delta}) | j \rangle \right|^2. \quad (5.5)$$

In practice, we can precalculate the factors of $\hat{U}_{\text{tg}}^\dagger | j \rangle$ to save matrix multiplications. An arbitrary time-evolution solver is needed to calculate $\hat{U}(\boldsymbol{\delta}) | j \rangle$, but such mathematics are well studied⁹⁶. In this case, we defer to pre-written libraries^{61,128}.

We must be careful that any comparison to the existing gate is fair. Many of the decoherence processes can be reduced by applying more laser power to perform the gate faster. This is easy for theorists to say, but impractical advice in reality; there is only so much power available. We need to impose a requirement that the peak power output of the driving field, $\max_t |f_n(t)|^2$, does not exceed the base gate.

This structure can be used to optimise the gate under any error model for the detunings. The work presented here focusses only on the two forms of qubit error, and leaves the motional error. The time-evolution operator for eq. (5.2) can be calculated with a terminating Magnus series (see eq. (2.16)), and multi-tone scheme has already been shown to be highly useful in these situations, in which the remnant undesired terms can be cancelled order by order^{42,110,132}. We will illustrate the method using normally distributed qubit-average and -splitting errors, as this is among the most likely model for errors with imperfect calibration.

5.2.1 Parametrisation

As in the optimisations in chapter 4, we seek some parametrisation that will allow us to use an unconstrained optimisation routine. The complex numbers $\{c_j\}$ are all very naturally parametrised by a real amplitude and phase. With arbitrary shapes,

the detuning also needs to vary, but it cannot do so freely without allowing the power to grow unbounded. Instead, we fix ϵ to fix the peak-power output at the maximal allowed value, reducing the number of parameters by one.

The scaling required for an arbitrary set of $\{c_j\}$ has no closed-form solution. It is possible in theory to scan through the duration of the gate and locate various maxima, but with the number of these extrema not known for any given parameter vector—there could be fewer than expected—it can be fiddly to ensure that the global maximum was found. A more elegant method is to recast the problem. All extrema coincide with the locations of the roots of the derivative of the power

$$\frac{\partial}{\partial t} |f_n(t)|^2 = \frac{\epsilon_n^2}{16} \sum_{j,k} c_{n,k} c_{n,j}^* e^{i(j-k)\epsilon_n t}. \quad (5.6)$$

The exponentials can be treated as powers of a variable $z = e^{i\epsilon_n t}$, and so multiplying through by z^{n-1} and relabelling indices leads to a polynomial

$$\left[\sum_{k=0}^{n-2} \left(\sum_{j=1}^{k+1} c_j c_{j-k+n-1}^* \right) (k-n+1) z^k \right] + \left[\sum_{k=n}^{2n-2} \left(\sum_{j=k-n+2}^n c_j c_{j-k+n-1}^* \right) (k-n+1) z^k \right] = 0. \quad (5.7)$$

All n complex roots of this z_ℓ can be found by well-established eigenvalue methods on the companion matrix of the polynomial⁹⁶. The roots are related to the temporal locations of the extrema $t_{\ell,m}$ by

$$\epsilon_n t_{\ell,m} = \arg(z_\ell) - i \ln|z_\ell| + 2\pi m. \quad (5.8)$$

The only roots of interest are in the first period and are real, so $m = 0$ and $|z_\ell| = 1$, and the peak power comes from testing the $2n - 2$ or fewer resulting abscissae. With this scaling in place, we have attained a smooth surjection from \mathbb{R}^ℓ onto the search space, and can now use the same BFGS method as in section 4.3.2.

5.2.2 Reduction of dimensionality

The state-space dimension for two qubits is four, and so the natural evaluation of the average gate fidelity requires numerically calculating four solutions to the Schrödinger equation. This can be reduced by more mathematically examining the effects of different detunings on the system. We first note that the multidimensional normal distribution centred on zero that we are using as an error model is completely symmetric under any sign flips; all choices of the signs in $\omega(\pm\delta_1, \pm\delta_2)$ give the same result. Further, the integral over all possible detunings in the figure of merit

TABLE 5.1—Frame transformations showing the equivalence of various combinations of errors and starting states. The second and third columns show pairs of basis states and error configurations that will exhibit identical dynamics to $|gg\rangle$ under a Hamiltonian with an average-frequency shift of δ_{avg} and a frequency split of δ_{spl} . Note that any single Pauli-Y flip exchanges δ_{avg} and δ_{spl} , and a flip of the first qubit introduces negatives on both.

$\hat{\mathcal{V}}$	$\hat{\mathcal{V}} gg\rangle$	$\hat{\mathcal{V}}\hat{\mathcal{H}}_{\text{MS}}(\delta_{\text{avg}}, \delta_{\text{spl}})\hat{\mathcal{V}}^\dagger$
$\hat{\mathbb{I}}$	$ gg\rangle$	$\hat{\mathcal{H}}_{\text{MS}}(\delta_{\text{avg}}, \delta_{\text{spl}})$
$\hat{\sigma}_y^{(1)}$	$-i eg\rangle$	$\hat{\mathcal{H}}_{\text{MS}}(-\delta_{\text{spl}}, -\delta_{\text{avg}})$
$\hat{\sigma}_y^{(2)}$	$-i ge\rangle$	$\hat{\mathcal{H}}_{\text{MS}}(\delta_{\text{spl}}, \delta_{\text{avg}})$
$\hat{\sigma}_y^{(1)} \otimes \hat{\sigma}_y^{(2)}$	$- ee\rangle$	$\hat{\mathcal{H}}_{\text{MS}}(-\delta_{\text{avg}}, -\delta_{\text{spl}})$

eq. (5.5) ensures that both positive and negative detunings will be accounted for, and have equal weight.

In the discussion surrounding fig. 5.1, we argued qualitatively that the effects of an average-frequency shift on the state $|gg\rangle$ would be similar to those of a splitting of the qubit frequencies on a singly excited state such as $|ge\rangle$. These effects are in fact quantitatively equal. Further, there is a symmetry within each pair of states, in that the $|gg\rangle \leftrightarrow |ee\rangle$ transition will have very closely related Bell-state-creation fidelities for both initial states in the basis.

Formally, we introduce an explicit parametrisation of the Hamiltonian eq. (5.2) as $\hat{\mathcal{H}}_{\text{MS}}(\delta_{\text{avg}}, \delta_{\text{spl}})$. It was shown in eq. (2.10) that a time-independent unitary frame transformation $\hat{\mathcal{V}}$ transforms the Hamiltonian $\hat{\mathcal{H}} \rightarrow \hat{\mathcal{H}}' = \hat{\mathcal{V}}\hat{\mathcal{H}}\hat{\mathcal{V}}^\dagger$, and similar for the time-evolution operator $\hat{\mathcal{U}}$. Now, consider the average gate fidelity calculated over the natural basis set $\{|gg\rangle, |ge\rangle, |eg\rangle, |ee\rangle\}$. The four possible combinations of applying or not applying the operators $\{\hat{\sigma}_y^{(1)}, \hat{\sigma}_y^{(2)}\}$ transform the state $|gg\rangle$ into one of the other four basis states, and the evolution is then equivalent to some exchange of the errors. All four equivalences are shown in table 5.1.

Under the assumption that the weight function is symmetrical around the zero point of each error individually, one only need evolve the states $|gg\rangle$ and $|ge\rangle$ to calculate the loss function eq. (5.5) exactly. The gate fidelity for the other two basis states averaged over the symmetric error probability function will be equal to these two. This halves the number of time evolutions that must be carried out, which is the core inner routine in the calculation. In general the weighted gate fidelity will differ between $|gg\rangle$ and $|ge\rangle$ if the standard deviations of the error models for the two parameters are different.

These similarities also show that a driving function optimised to reduce the effects of average-frequency shifts for all starting states will also be optimised for non-degenerate qubit levels. This does not extend as far as motional detunings, however. There is no similarity transformation between the qubit basis states that relates one of the other errors to a motional error.

5.2.3 Quadrature

The only remnant issue in the optimisation is the evaluation of the integral over all possible detunings in eq. (5.5). The time integrations of the Hamiltonian necessary to calculate $\hat{U}(\delta)$ are expensive, and it is desirable to reduce the number of evaluations of the integrand. Numerically this becomes trickier when the integral is multidimensional, as in this case.

Standard one-dimensional integration is often carried out by approximating the integrand in the integration region by a low-order polynomial. The calculation is performed as

$$\int_a^b dx f(x) \approx \sum_{j=1}^n w_j f(x_j) \quad \text{for } x_j = a + (j-1) \frac{b-a}{n-1}, \quad (5.9)$$

for some weights $\{w_j\}$ that are independent of the integrand. They are found by defining a related function $f'(x) = \sum_j f(x_j) \ell_j(x)$ that is the Lagrange interpolation of f , where the $\{\ell_j\}$ are $(n-1)^{\text{th}}$ -degree polynomials that satisfy $\ell_j(x_k) = \delta_{jk}$ for abscissae on the grid. These are given by $\ell_j(x) = \prod_{k \neq j} (x - x_k) / (x_j - x_k)$, with the numerator providing a multiplicative factor of zero for grid points other than x_j and the scaling fixed to ensure $\ell_j(x_j) = 1$. The function f' can then be integrated analytically, since the $f(x_j)$ are known; the weights in eq. (5.9) are $w_j = \int_a^b dx \ell_j(x)$.

This is a description of the *Newton–Cotes* family of quadrature rules, of which the trapezium rule (linear) and Simpson’s rule (quadratic) are the most well known. It is not usually worth moving to higher degrees of polynomial interpolation, which require more evaluations of the integrand but do not typically significantly improve error rates for non-polynomial functions. Instead, one may use the trapezium rule and subdivide the region recursively until some tolerance threshold is reached. This immediately poses two problems: the region here is infinite; and the number of function evaluations is not predictable, making it more difficult to allocate computational resources.

Sticking, for now, to one dimension, it is possible to solve both of these problems, and achieve higher degree rules with the same number of points. Newton–Cotes rules require n function evaluations on an evenly spaced grid to reach an interpolation with a polynomial of degree $n-1$. However, there are more degrees of freedom available to construct higher-order rules. We previously assumed that the grid points $\{x_j\}$ would be evenly spaced, but this is not necessary. In fact, by a suitable choice of locations and quadrature weights, we can use only n function evaluations to achieve a degree of $2n-1$ and account for a weight function within

TABLE 5.2—Commonly used weight functions and associated regions of integration used with Gaussian quadrature⁹⁶. The regions and natural weight functions are related, but analytic variable transformations in the integrands can convert any finite region to any other; the *form* of the weight is more important than the scaling.

Weighting $\omega(x)$	Region Ω	Polynomial set
1	$[-1, 1]$	Legendre
$1/\sqrt{1+x^2}$	$[-1, 1]$	Chebyshev
$(1-x)^\alpha(1+x)^\beta$	$[-1, 1]$	Jacobi
$x^\alpha \exp(-x)$	$[0, \infty)$	Laguerre
$\exp(-x^2)$	$(-\infty, \infty)$	Hermite

the integrand itself. This is *Gaussian quadrature*.

We begin by considering the set of smooth square-integrable real-valued functions as a Hilbert space under a family of inner products defined by

$$\langle f, g \rangle_{\omega, \Omega} = \int_{\Omega} dx \, \omega(x) f(x) g(x), \quad (5.10)$$

for some region Ω and a weighting function $\omega(x)$ that is positive everywhere. One can construct a series of *orthogonal polynomials* $\{P_k\}$, defined by $\langle P_n, P_m \rangle \propto \delta_{nm}$, by starting from the basis of monomials $\{x^k \mid k \in \mathbb{N}^0\}$ and applying the Gram-Schmidt process. The most commonly used weight functions and regions produce well-known sets of orthogonal polynomials, tabulated in table 5.2.

Using polynomial division, an arbitrary function $f(x)$ in the vector space can be rewritten in terms of a quotient q and remainder r as

$$f(x) = q(x)P_m(x) + r(x), \quad \text{where } \text{degree}(r) < m. \quad (5.11)$$

Since the set of orthogonal polynomials spans the function space, we can decompose $q(x) = \sum_k c_k P_k(x)$ for some constants $\{c_k\}$. Let us now assume that f is a polynomial function of degree strictly less than $2m$. In this case q must be of a degree less than m , and the integral we are interested in is reduced by orthogonality:

$$\int_{\Omega} dx \, \omega(x) f(x) = \left(\sum_{k < m} c_k \int_{\Omega} dx \, \omega(x) P_k(x) P_m(x) \right) + \int_{\Omega} dx \, \omega(x) r(x), \quad (5.12)$$

and all terms in the summation are zero. For quadrature we discretise to find

$$= \sum_{j=1}^m w_j q(x_j) P_m(x_j) + \sum_{j=1}^m w_j r(x_j), \quad (5.13)$$

with equality due to the prior polynomial assumption of f . The first summation can be made zero as required by choosing the sample locations $\{x_j\}$ to be the roots

of the polynomial P_m . This leaves only a term of degree less than m , and the same Laguerre-interpolation methods used to derive Newton–Cotes rules will find the necessary values of the $\{w_j\}$ to make the discretisation exact. Both the abscissae and weights are independent of the integrand—other than the fixed weight function in the inner product—and can be calculated ahead of time, or, more realistically, looked up in reference works^{26,27,96}.

The weight functions we are concerned with here are all of the form $\exp(-x^2)$ for some coordinate scaling by the standard deviation in the error model. Using the Gaussian quadrature rules we have just described over the set of *Hermite polynomials* allows us to calculate each separate integral over the infinite integral to high precision with few evaluations of the integrand. This is because the detunings are relatively small, and $\hat{\mathcal{U}}(\delta)$ varies sufficiently slowly—in a mathematical, not physical, sense—with respect to drifts in the detunings. We could now build up a multidimensional rule by iterating the integrations, which would require m^d evaluations to interpolate at degree m over d dimensions.

While this achieves better results with fewer evaluations than adaptive-width trapezium rules, in certain circumstances it can be possible to go even further. For such symmetric weight functions as $\exp(-\mathbf{r} \cdot \mathbf{r})$, there is a body of literature that finds specific rules for multidimensional integrals^{26,121}. Orthogonal polynomials do not easily generalise to higher-dimensional spaces, and so such rules are rather less systematic, and more scattershot in the available degrees and required number of evaluations. Still, for the cases of two-dimensional integration with a Gaussian weight function, we will use rules from Cools²⁶ for degrees that have known more efficient solutions, and fall back on iterated Gauss–Hermite quadrature otherwise.

5.3 Results

In principle, one would set the standard deviations in the error model by measuring drift over the system, and producing an estimated error model. With no particular experimental setup in mind, we instead run the optimisations over several different values of these *hyper-parameters* to investigate the families of gate produced. As expected from the analysis of section 5.2.2, optimising the gate over all states resulted in the same solutions, no matter whether the error model accounted for independent average offset and qubit splittings or only one of these. The symmetry under exchange of states ensured that all effects invariably contributed to the figure of merit. More families of solutions were found when integrating over only a single dimension, simply because of the reduced computational cost, and all

Tones	τ_n/τ_1	δf_n	ϵ_n/ϵ_1		$c_{n,1}$	$c_{n,2}$	$c_{n,3}$	$c_{n,4}$	$c_{n,5}$	$c_{n,6}$
2	3.368	0.132	0.297	$ c $ ϕ/π	0.066 -0.032	0.934 0				
3	3.185	0.033	0.314	$ c $ ϕ/π	0.103 -0.005	0.979 -0.003	0.090 0			
4	4.836	0.555	0.207	$ c $ ϕ/π	0.051 -0.609	0.405 -0.817	0.539 0.108	0.359 0		
5	4.542	0.622	0.220	$ c $ ϕ/π	0.048 -0.899	0.450 -0.930	0.516 0.045	0.414 -0.242	0.183 0	
6	6.529	0.482	0.153	$ c $ ϕ/π	0.055 -0.616	0.098 -0.785	0.413 -0.954	0.733 0.007	0.215 -0.043	0.128 0

TABLE 5.3—Multi-tone driving schemes of the Mølmer-Sørensen gate that are optimised to reduce the effects of normally distributed static qubit frequency errors. The fields are as described by eq. (5.4), where the complex amplitudes $c_{n,k} = |c_{n,k}|e^{i\phi_{n,k}}$. These amplitudes are scaled such that $c_{1,1} = 1$ retrieves the standard single-tone dynamics. The single-tone gate has a constant total amplitude of 1, whereas the multi-tone schemes vary by up to δf_n over the course of the gate. The phase of each most-detuned tone is arbitrarily chosen to be zero.


of these solutions remained valid and stable when used as the starting point for optimisations over both errors simultaneously.

It is worth highlighting that in multidimensional optimisations such as these, the convergence of the algorithm on a particular solution does not require it to be globally optimal, only locally. This can result in schemes that have the best response to errors with respect to small changes in the control parameters, but a much larger modification could produce significantly better results. The best general schemes for numbers of tones between two and six are shown in table 5.3.

Unsurprisingly, one stable family of gates across a wide variety of error spreads is a minor perturbation to a single-tone gate. In these, one tone commands around 90% of the total power, making it roughly equivalent to the base gate, but potentially drawing out more loops in phase space before the completion of the gate, depending on which of the tones has the most power. For two- and three-tone gates, this style was the only stable solution that produced better fidelities than the base gate for a fixed peak power; for various fixed detunings, these schemes exhibited approximately a three-times reduction of infidelity. It is likely that at this level, there were simply not enough degrees of freedom for the optimiser to produce any more meaningful improvements.

This story changed completely once the fourth tone was added, however. These schemes diverged from simple perturbations of the single-tone gate. Figure 5.2 shows the gate infidelity for the different schemes in table 5.3 at fixed values of

FIGURE 5.2—Gate infidelities for the optimised driving schemes with the best overall performance compared to the standard Mølmer–Sørensen scheme. For the sake of illustration, the gate infidelities are evaluated here for detunings in the ratio $\delta_{\text{avg}} = 2\delta_{\text{spl}}$, but the behaviour is qualitatively the same for any ratio. An error which causes the single-tone gate to leave the error-correction threshold of 99.9% causes an infidelity of only 2.5×10^{-5} when four or more tones are used. The two- and three-tone gates are minor modifications of the standard driving, yet produce a three- to four-times improvement over the range of meaningful infidelities. The inset shows the amplitude variation of each scheme over the gate duration.



the detunings. The response with respect to the errors is qualitatively the same no matter what particular pair of the average and splitting errors are chosen. In this figure, we arbitrarily evaluated the gate infidelity for detunings in a ratio $\delta_{\text{avg}} = 2\delta_{\text{spl}}$ for the sake of making a one-dimensional figure to better illustrate the scaling behaviours. The two- and three-tone gates did not achieve better than the δ^2 scaling of the base gate. All schemes with four or more tones achieved a scaling of δ^4 in the regions of greatest error.

This improvement in scaling brings order-of-magnitude improvements over the single-tone gate. At two of the points often cited as the thresholds for fault-tolerant quantum computing, 99% and 99.9%⁶⁴, the perfectly implemented four-tone gate would instead have infidelities of 10^{-3} and 2.5×10^{-5} respectively, improvements of approximately 10- and 40-times. Alternatively, one could ask *how much larger can the error be while maintaining fault tolerance?* The same two thresholds are breached for the four-tone gate with errors 1.8 and 3.2 times larger—a sizeable improvement. The story is largely the same for the five- and six-tone gates. While the constant factors of the infidelity can be reduced with the extra tones, there remains insufficient degrees of freedom to achieve greater scaling.

We must also address the elephant in the room, that is the failure of the infidelities to go to zero at zero error, and indeed how the standard gate beats the “optimised” gates for certain lower errors. Both of these are explained by the numerical target. There is no requirement built-in to the figure of merit that the gate is perfect under ideal conditions, only that the expectation of the infidelity over the distribution of errors is minimal. With the standard deviation hyper-parameters set sufficiently large as to involve errors that prevent fault tolerance, the contribution of the smallest detunings is minor. The expectation is dominated by other regions of

possible errors.

The flattening out of the infidelities of the optimised schemes may well not be a fundamental issue. Beyond a certain limit, the numerical uncertainty from the quadrature or the time evolution may make it impossible for the optimiser to make further progress. If so, it is likely that the differences in the parameters between the schemes as presented in table 5.3 and hypothetical schemes that maintain the δ^4 scaling all the way to zero error would be sufficiently small as to be indistinguishable in experimental realisations.

A notable feature of fig. 5.2 is that only even-numbered tones appear to significantly affect the response of the gate. The infidelity of the three-tone gate looks almost identical to that of the two-tone gate, and similar between the four- and five-tone gates. The inset in the figure shows how the amplitude varies across the gate time, but this does not explain the similarities seen in the infidelity responses. In particular, the four- and five-tone gates do not seem to have too much in common in this respect.


A better tool to compare the effects of different schemes is phase-space analysis. The two terms of the Magnus expansion of the ideal Mølmer–Sørensen Hamiltonian as originally given in eq. (3.27) are

$$\begin{aligned}\hat{M}_1(t) &= -i\frac{\eta\Omega}{2}\hat{S}_y \int_0^t dt_1 (e^{-i\epsilon t_1}\hat{a}^\dagger + e^{i\epsilon t_1}\hat{a}), \text{ and} \\ \hat{M}_2(t) &= i\frac{(\eta\Omega)^2}{4}\hat{S}_y^2 \int_0^t dt_1 \int_0^{t_1} dt_2 \sin(\epsilon(t_2 - t_1)).\end{aligned}\tag{5.14}$$

Considering the two non-zero eigenstates of the \hat{S}_y operator with eigenvalues ± 1 , the first term produces an effective Hamiltonian $\hat{\mathcal{H}}_{\text{eff}}(t) = \hat{\mathcal{D}}\left(\pm \frac{i\eta\Omega}{2\epsilon}(e^{-i\epsilon t} - 1)\right)$, while the second is the two-qubit entangling interaction.

This displacement term permits a geometrical interpretation of the effects of the gate. The two eigenstates gain a phase relative to each via their motional displacement, shown by the area enclosed by the displacement: it is directly proportional to the angle of the applied \hat{S}_y^2 rotation. This area has to correspond to an angle $(4n + 1)\pi/4$ for some integer n to have the intended entangling effect. The displacement must return to zero at the completion of the gate to avoid unwanted entanglement between the qubits and the motion. The variation of the displacement of the motion over the course of the gate offers the easiest insight into the similarities and differences of driving schemes, as it encodes information on both the status of the motional mode, and the angle of the desired \hat{S}_y^2 rotation simultaneously.

FIGURE 5.3—Motional phase-space trajectories of the different multi-tone gates also plotted in fig. 5.2, with the same peak power usage and different gate times. Colour variation shows relative time through the gate, starting from dark purple. Only even numbers of tones cause structural changes to the paths; this matches the fidelity responses seen in fig. 5.2.



Phase-space expectation values for the positive eigenstate of \hat{S}_y are plotted in fig. 5.3 for the base Mølmer–Sørensen scheme, and each of the other five gates presented in table 5.3. The two non-dimensionalised operators $\hat{x} \propto \hat{a}^\dagger + \hat{a}$ and $\hat{p} \propto \hat{a}^\dagger - \hat{a}$ correspond to the real and imaginary components of the displacement respectively, which draws out a circle when driving at constant power. This particular shape has long been known to be unnecessary, however, provided it encloses the correct area^{79,116}.

The similarities between the infidelity responses shown in fig. 5.2 of odd–even pairs of driving fields are explained by fig. 5.3. The phase-space trajectory ultimately determines the effects of detunings on the fidelity of the resulting gate. While the amplitude variations of each pair do not seem entirely linked, the addition of the extra tone does not appear to offer enough freedom to structurally change the path. Without analytic access to the propagators, it is difficult to explain exactly why this should be the case. One thought might be that there is some particular symmetry of the problem that discourages odd tones, but if this were the case, we should generally expect that odd-numbered tones—except for the base-gate case—should have negligible amplitudes in order to allocate the power elsewhere. The tabulated values in table 5.3 do not show any such pattern, nor really any other relation.

Throughout this chapter we have only considered detunings of the qubit frequencies. Notably, we did not attempt to minimise errors against both these frequencies and the motional frequency simultaneously with the same parametrisation. Motional frequencies alone have been considered with this same parametrisation^{42,110,132}, returning a distinct cardioid shape of the phase-space path.

The methods used in this chapter reproduce these prior results when handling only a motional uncertainty, but it is more interesting to attempt to handle all three errors at once. Unfortunately, we found no suitable schemes with this parametrisation alone. This can be understood through the shapes of the phase-space paths. It can be shown analytically⁴² that keeping the average motional displacement at

the zero point, and minimising its maximum distance from the origin produces gates that are most resilient against mis-sets in the motional frequency. Physically, this seems natural: the greater the intended phase-space displacement, the greater the error rate caused by slight mis-sets in the parameters of the motion. Similarly, keeping the average displacement centred at zero keeps any residual qubit–motion entanglement small in magnitude, since it is never large at any point. With this in mind, it becomes clear why the same parametrisation cannot handle qubit and motional noise simultaneously; the centres of the displacements in fig. 5.3 is decidedly non-zero, and the maximum displacement generally grows beyond the base-gate case as more tones are added.

Still, it is worth highlighting that this parametrisation offers significant improvements in qubit errors when motional errors can be tightly controlled. This may not often be the case, but in a system plagued mostly by stray magnetic fields, for example, one can use the smooth control present in this chapter to achieve ten- to forty-times improvements at critical points of frequency drift. Above three tones, the Fourier-series parametrisation showed quadratic improvements in the *scaling* of the infidelity with respect to error, suggesting that some fundamental term can be entirely nulled with this method.

5.4 Outlook

The shaped pulses described in table 5.3 showed significant improvements in fidelities for unknown qubit-frequency errors, though their inability to simultaneously handle errors in the motional frequency hampers their applicability in the real world. Further, control schemes based on precise, continuous amplitude shaping can suffer from increased calibration requirements in some experimental set-ups. Some methods of control-field synthesis exhibit non-linear responses to amplitude modulation, which can make implementation of these sequences difficult¹³². This is not to say that the Mølmer–Sørensen gate has no future—it is still commonly used by the largest ion-trap-based quantum-computing companies IonQ¹⁴ and Honeywell⁹³—only that this particular parametrisation of the control field may not be a silver bullet for error mitigation on its own.

There have been other methods of gate shaping that have shown promise. One can maintain a constant amplitude throughout the gate, and shape the phase-space trajectories by introducing discrete jumps in the applied phases of the fields⁸⁰. Alternatively, one can use single-qubit spin-echo techniques to mitigate some of the decoherence effects⁷⁶. The qubit error term $\hat{\sigma}_z$ does not commute with the gate

Hamiltonian, preventing the spin-echo method from achieving perfect suppression of these errors, but it can still improve fidelities.

The Fourier-series amplitude-shaping method presented in this chapter appears to be more successful at reducing motional decoherence than qubit miscalibrations^{110,132}, and so mixing different parametrisation schemes may well be the best path forwards, with qubit-error considerations best solved by other means. Very recent work has combined multi-tone methods with spin-echo techniques, achieving good isolation against both motional- and qubit-frequency errors⁷⁷. With these other methods seemingly yielding better results and no immediate experimental collaborations, we did not take the gate parametrisation of this chapter further. Instead, we now turn to look at a previously more fundamental limitation of trapped-ion gates: the linearity approximation in their Hamiltonian.

Chapter 6

Entangling Gates With Strong Coupling

COAUTHORSHIP

The initial mathematical expansion of the entangling interaction in this chapter was proposed by Mahdi Sameti, and was subsequently polished into the complete form in collaboration with me and Florian Mintert. I performed all the numerical simulations of the gates, and wrote the algorithms to programmatically evaluate the functional constraints on the driving profiles in symbolic form. Mahdi and I worked together to find particular solutions to the resulting systems. The work in this chapter was also published in ref. 104.

The very first step in deriving the interaction Hamiltonian for a driving field with trapped ions is to make a linearity approximation. This is the Lamb–Dicke approximation that we introduced in section 3.3.2, and have been using ever since. Physically, it is a requirement that the coupling between the qubit states and the shared motion of the ions is weak. Mathematically, we encoded it as $(n + 1)\eta^2 \ll 1$ for the Lamb–Dicke parameter η and motional occupation n . It limits the available interactions to only the first-order sidebands and the carrier, and is necessary for the original Mølmer–Sørensen interaction with constant power to be independent of the motional state.

With modern hardware, however, the necessity of maintaining this regime is becoming a more onerous burden. The highest fidelity gates must account for the breakdown of the approximation in their error budgets, where it is a non-negligible contribution^{4,40}. A common impediment to large-scale ion-trap quantum computing is the speed of gate operations. In order to increase speed, one must have more power available and consequently drive the gate with a larger value of the Lamb–Dicke parameter. As this increases, though, so do the non-linearities associated with the breakdown of the Lamb–Dicke regime, until they dominate all errors and reduce the fidelity well beyond any acceptable limit¹⁰⁷. These attempts to increase the speed have been accompanied by some numerical work to curtail these effects, but prior to our work, there had been nothing systematic nor analytic.

The non-linearities are also worsened by heating of the motional modes. With even modest values of the Lamb–Dicke parameter, if the motion is sufficiently excited, one will still see residual qubit–motion entanglement at the gate-completion time, and the entangling interaction itself will see its strength become phonon-number dependent. As the sizes of ion traps increase, so too does the heating rate on the ions, in general, meaning greater deviation from ideal gate operation.

For trapped-ion quantum processors to continue to scale up, without unrealistic requirements on heating control and laser power, it is necessary to be able to break through the linearity approximation. This chapter presents a functional approach to bringing trapped-ion gates outside the Lamb–Dicke approximation, by driving higher-order sidebands to null undesired terms at the gate time. We focus on the Mølmer–Sørensen interaction, though the essential techniques are just as applicable to $\hat{\sigma}_z \otimes \hat{\sigma}_z$ -type gates¹⁰¹. The recipe systematically removes terms of higher powers of the Lamb–Dicke parameter and the motional occupation as more sidebands are considered. Importantly, it does not require specific forms for the driving profile of each sideband, allowing it to be mixed with extant methods of suppressing other errors, such as the multi-tone gates discussed in the previous chapter.

6.1 Non-linear ion–motion interactions

The laser–ion interaction Hamiltonian for a single motional mode is given in eq. (3.19). In the weak-coupling regime, one expands the motional exponential as

$$\hat{D}(i\eta e^{i\omega_z t} \hat{a}^\dagger) = 1 + i\eta e^{-i\omega_z t} \hat{a} + i\eta e^{i\omega_z t} \hat{a}^\dagger + \mathcal{O}(\eta^2), \quad (6.1)$$


and neglects the higher-order terms. More precisely, one can use the Baker–Campbell–Hausdorff-derived eq. (3.22) to separate out the exponential into a sum

$$\begin{aligned} &= e^{-\frac{\eta^2}{2}} \sum_{j=0}^{\infty} \sum_{k=0}^{\infty} \frac{(i\eta)^{j+k}}{j!k!} e^{i(j-k)\omega_z t} \hat{a}^{\dagger j} \hat{a}^k \\ &= e^{-\frac{\eta^2}{2}} \sum_{k=-\infty}^{\infty} e^{ik\omega_z t} \hat{A}_k(\eta). \end{aligned} \quad (6.2)$$

The operators $\{\hat{A}_k\}$ change the motional state by k phonons, and are explicitly

$$\hat{A}_k(\eta) = \sum_{n=0}^{\infty} (i\eta)^{|k|} \frac{(-\eta^2)^n}{(n+k)!n!} \hat{a}^{\dagger n+k} \hat{a}^n \quad \text{for } k \geq 0, \text{ and } \hat{A}_{-k} = (-1)^k \hat{A}_k^\dagger. \quad (6.3)$$

FIGURE 6.1—Energy-level diagram for two ions and one motional mode driven by the Hamiltonian in eq. (6.6). Applying a drive on equal-order blue- and red-sideband transitions simultaneously, close to resonance, results in an effective qubit–qubit interaction with phonon-dependent strength Γ . If only the first-order sidebands are driven inside the Lamb–Dicke regime, this coupling is independent of the motion.



The complete interaction including the qubit-excitation operator $\hat{S}_+ = \sum_j \hat{\sigma}_+^{(j)}$ is

$$\hat{\mathcal{H}} = f(t) e^{-\frac{\eta^2}{2}} \sum_{k=-\infty}^{\infty} e^{ik\omega_z t} \hat{S}_+ \hat{A}_k + \text{H.c.} \quad (6.4)$$

Each of the different k corresponds to a well-defined transition as previously discussed: $k = 0$ is the carrier, positive values of k are the k th-order blue sideband, and negative values of k are the corresponding red sideband. As before, each transition has an associated frequency $k\omega_z$, corresponding to the total frequency of phonons that need to be added—or removed, in the case of the red-sideband transitions—to the system. For the most part, the only transitions that have any appreciable effect on the dynamics are ones which share their frequency with a component of the driving field.

With eq. (6.4), one can drive a series of two-photon processes that preserve the phonon count by targeting the equal-order red and blue sidebands with related drives. In particular, let us consider a driving field

$$f(t) = -i \frac{e^{\eta^2/2}}{\eta} \sum_{k>0} \left(f_k(t) e^{-ik\omega_z t} + (-1)^k f_k^*(t) e^{ik\omega_z t} \right). \quad (6.5)$$

The remaining time dependences of the $\{f_k\}$ are slow compared to the trap frequency ω_z so that each term only drives its closest sideband. The prefactor of $e^{\eta^2/2}$ cancels out its corresponding term in the Hamiltonian, while the initial phase factor $-i$ and the relative phase factors $(-1)^k$ are chosen such that the Hamiltonian, neglecting far-off-resonant processes, simplifies to

$$\hat{\mathcal{H}} = \frac{1}{\eta} \hat{S}_y \sum_{k>0} \left(f_k \hat{A}_k + f_k^* \hat{A}_k^\dagger \right). \quad (6.6)$$

This does not include any carrier transitions; we will not use them in the following work. The factor of $1/\eta$ is crucial to the rest of this chapter, but we need to advance

further in the derivations to fully motivate its inclusion.

Figure 6.1 illustrates the effects of the Hamiltonian eq. (6.6). If any of the f_k are non-zero, a two-photon two-qubit-entangling process is present with a phonon-dependent coupling strength Γ . In the special case of the first sideband and the Lamb–Dicke regime, this coupling strength is independent of the number of motional excitations, and the only other process is the qubit–motion interaction we have already discussed in detail.

The standard linear Mølmer–Sørensen interaction is retrieved from eq. (6.6) by addressing only the first-order sideband transitions, *i.e.* driving only f_1 . The original formulation of the gate corresponds to $f_1(t) = \frac{\epsilon}{4}e^{i\epsilon t}$. We have already seen the Magnus-expansion approach to solving the propagators in both eqs. (3.27) and (5.14), which includes a term $[\hat{A}_1, \hat{A}_{-1}] = \eta^2 - 2\eta^4\hat{a}^\dagger\hat{a} + \mathcal{O}(\eta^6)$. In the Lamb–Dicke limit, this is truncated to its leading order, and the expansion terminates.

To go further, let us define a shorthand notation for iterated time integration as

$$\{f\} = \int_0^t dt_1 f(t_1) \quad \text{and} \quad \{f\{g\}\} = \int_0^t dt_1 f(t_1) \int_0^{t_1} dt_2 g(t_2), \quad (6.7)$$

and so forth for further nesting. In this form, the time-evolution operator for the base Mølmer–Sørensen gate, accounting for only the lowest orders of η , is*

$$\hat{\mathcal{U}}_0(t) = \exp\left(\hat{S}_y(\{f_1\}\hat{a}^\dagger - \{f_1^*\}\hat{a}) + i\hat{S}_y^2 \text{Im}\{f_1\{f_1^*\}\}\right). \quad (6.8)$$

This operator has three processes: the desired \hat{S}_y^2 , the unwanted $\hat{S}_y\hat{a}^\dagger$, and its reverse counterpart $\hat{S}_y\hat{a}$. Achieving an ideal gate is tantamount to solving the coupled conditions $\{f_1\} = 0$ and $\text{Im}\{f_1\{f_1^*\}\} = \theta$, for a desired entangling angle θ . Bell-state creation corresponds to $\theta = \pi/8$.†

We have now reached our departure point from known theory, and from the Lamb–Dicke regime. We now must consider Hamiltonians that are *non-linear* in the motional operator \hat{a} , and contain higher-order terms in η . In principle, our approach will be to find a series expansion for the time-evolution operator akin to eq. (6.8), and then apply the constraints that all processes except for \hat{S}_y^2 go to zero at the gate time.

We will seek to nullify unwanted terms order-by-order in η by constructing a product expansion as

$$\hat{\mathcal{U}} = \hat{\mathcal{U}}_0\hat{\mathcal{U}}_1 \cdots \hat{\mathcal{U}}_d\hat{\mathcal{U}}_{\text{error}}, \quad (6.9)$$

*To reduce notational noise, this chapter uses the convention $\hbar = 1$.

† $\hat{S}_y^2 = 2 + 2\hat{\sigma}_y \otimes \hat{\sigma}_y$, hence the factor-of-two difference from the angle that might be expected.

where the final term is never explicitly calculated, but is guaranteed to only contain higher-order terms. The expansion is built by a series of frame transformations, each of which removes the terms of the lowest order of η . Starting from the base Hamiltonian, which we shall now label $\hat{\mathcal{H}}_0$, we move to the next Hamiltonian by the transformation

$$\hat{\mathcal{H}}_{j+1} = \hat{\mathcal{U}}_j^\dagger \hat{\mathcal{H}}_j \hat{\mathcal{U}}_j + i(\partial_t \hat{\mathcal{U}}_j^\dagger) \hat{\mathcal{U}}_j. \quad (6.10)$$

The first unitary is exactly $\hat{\mathcal{U}}_0$ as described by eq. (6.8). To remove the terms lowest-order in η from the Hamiltonians, we define all subsequent unitaries as

$$\hat{\mathcal{U}}_j = \exp\left(-i\{\hat{\mathcal{H}}'_j\}\right) \quad \text{for } j > 0, \quad (6.11)$$

and $\hat{\mathcal{H}}'_j$ is the terms of order exactly η^j from $\hat{\mathcal{H}}_j$. Note that $\hat{\mathcal{U}}_j$ is not necessarily a solution to any Schrödinger equation; this is not required for the expansion.

We do not need to directly calculate any matrix exponentials to evaluate eq. (6.10). The first term permits a direct Baker–Campbell–Hausdorff expansion—using the result of eq. (2.13)—yielding

$$\hat{\mathcal{U}}_j^\dagger \hat{\mathcal{H}}_j \hat{\mathcal{U}}_j = \frac{1}{0!} \hat{\mathcal{H}}_j + \frac{1}{1!} \left[i\{\hat{\mathcal{H}}'_j\}, \hat{\mathcal{H}}_j \right] + \frac{1}{2!} \left[i\{\hat{\mathcal{H}}'_j\}, \left[i\{\hat{\mathcal{H}}'_j\}, \hat{\mathcal{H}}_j \right] \right] + \cdots. \quad (6.12)$$

The second term requires an evaluation of the derivative of a matrix exponent. This derivative for a general matrix $\exp \hat{X}$ is⁴⁶

$$\partial_t e^{\hat{X}} = \left(\int_0^1 d\alpha e^{\alpha \hat{X}} (\partial_t \hat{X}) e^{-\alpha \hat{X}} \right) e^{\hat{X}}, \quad (6.13)$$

which qualitatively is the continuous application of the chain and product rules of differentiation. Expanding the integrand into a sum with eq. (2.13) and directly integrating the resulting power series in α gives the relation

$$\left(\partial_t e^{\hat{X}} \right) e^{-\hat{X}} = \frac{1}{1!} \partial_t \hat{X} + \frac{1}{2!} \left[\hat{X}, \partial_t \hat{X} \right] + \frac{1}{3!} \left[\hat{X}, \left[\hat{X}, \partial_t \hat{X} \right] \right] + \cdots, \quad (6.14)$$

which relates to the $(\partial_t \hat{\mathcal{U}}_j^\dagger) \hat{\mathcal{U}}_j$ term by $\hat{X} = i\{\hat{\mathcal{H}}'_j\}$. The numeric factors in eqs. (6.12) and (6.14) differ slightly for d nested commutators due to the integration of a factor of α^d . The derivative of the time-integrated term $\{\hat{\mathcal{H}}'_j\}$ is exactly $\hat{\mathcal{H}}'_j$ by the standard rules of anti-derivatives. Altogether, this leads to a series form of eq. (6.10) as

$$\hat{\mathcal{H}}_{j+1} = \sum_{m=0}^{\infty} \frac{1}{m!} \underbrace{\left[i\{\hat{\mathcal{H}}'_j\}, \cdots \left[i\{\hat{\mathcal{H}}'_j\}, \hat{\mathcal{H}}_j - \frac{1}{m+1} \hat{\mathcal{H}}'_j \right] \cdots \right]}_{m \text{ commutators}}. \quad (6.15)$$

Since $\hat{\mathcal{H}}'_j$ is all the order- η^j terms from $\hat{\mathcal{H}}_j$, the first term of the summation only contains terms of order η^{j+1} and greater, while the subsequent elements have terms of order $\eta^{(m+1)j}$ and greater.

We have now arrived at a perturbative method for approximating the true time-evolution operator for the non-linear Mølmer-Sørensen-like Hamiltonian of eq. (6.6). The result is a propagator

$$\hat{\mathcal{U}} = \hat{\mathcal{U}}_0 \hat{\mathcal{U}}_1 \cdots \hat{\mathcal{U}}_d \hat{\mathcal{U}}_{\text{error}}, \quad \text{where } \hat{\mathcal{U}}_j = \exp\left(-i\{\hat{\mathcal{H}}'_j\}\right) \text{ for } j > 0, \quad (6.16)$$

and the $\hat{\mathcal{H}}'_j \propto \eta^j$ are the leading-order terms of the Hamiltonians from eq. (6.15). The error term is never calculated exactly, but if the expansion is made up to $\hat{\mathcal{U}}_d$, then the generator of the error term contains only terms of higher order than η^d . Further, we have constructed the terms without reference to the precise forms of the driving functions of each sideband.

We require that $\hat{\mathcal{U}} = \hat{\mathcal{U}}_{\text{MS}} \hat{\mathcal{U}}_{\text{error}}$ to make a valid gate, which imposes a series of conditions. The desired dynamics of $\hat{\mathcal{U}}_{\text{MS}}$ can be achieved if all the generators of the $\hat{\mathcal{U}}_d$ are zero, except for the terms proportional to the operator \hat{S}_y^2 with no motional dependence. We expand

$$\hat{\mathcal{H}}'_j = \eta^j \psi_j[f] \hat{S}_y^2 + \eta^j \sum_{\hat{X}} \chi_{j,\hat{X}}[f] \hat{X} \quad (6.17)$$

for some scalar functionals ψ and χ of the vector of sideband driving profiles f , and some operators that are calculated by eq. (6.15). Each operator \hat{X} must individually be cancelled out, and if so, the remaining \hat{S}_y^2 terms can all be brought into one exponential, leaving the final conditions

$$\sum_j \eta^j \psi_j[f] = \frac{\pi}{8} \quad \text{and} \quad \chi_{j,\hat{X}}[f] = 0 \text{ for all } j \text{ and } \hat{X}. \quad (6.18)$$

Assuming that the driving profiles are implemented perfectly such that the gate operator is $\hat{\mathcal{U}}_{\text{MS}} \hat{\mathcal{U}}_{\text{error}}$, we can calculate the minimal dependence on η . An expansion up to the term $\hat{\mathcal{U}}_{d-1}$ has an error term of the form

$$\hat{\mathcal{U}}_{\text{error}} = \exp\left(-i\eta^d \hat{R}\right) \quad (6.19)$$

for some unknown Hermitian \hat{R} that can itself contain factors of η . The gate

infidelity over a basis $|k\rangle$ is

$$\begin{aligned}
 I &= 1 - \sum_k \left| \langle k | \hat{\mathcal{U}}_{\text{error}} | k \rangle \right|^2 \\
 &= 1 - \sum_k \left\langle k \left| \left(1 - i\eta^d \hat{R} + \mathcal{O}(\eta^{2d} \hat{R}^2) \right) \right| k \right\rangle \left\langle k \left| \left(1 + i\eta^d \hat{R} + \mathcal{O}(\eta^{2d} \hat{R}^2) \right) \right| k \right\rangle \quad (6.20) \\
 &= \mathcal{O}(\eta^{2d}),
 \end{aligned}$$

since all powers of \hat{R} remain Hermitian. In other words, each additional step of the expansion provides the constraints needed to improve the error scaling with respect to the Lamb–Dicke parameter by a factor of η^2 . The base gate accounts for all terms of order η but not η^2 , so its infidelity scaling is $\mathcal{O}(\eta^4)$.

We may now discuss the factor of $1/\eta$ that we identified in eq. (6.5). At all stages of the perturbative expansion, it was crucial that we could identify the terms of lowest order in η . The general solutions to the Mølmer–Sørensen interaction all require a driving field that scales in power as $1/\eta$, which corresponds physically to matching the coupling strength of the first-order sideband. If this is left in the individual sideband profiles, there is a potential factor of η^{-1} lurking that prevents eq. (6.15) from strictly increasing the powers of η . By extracting the negative power immediately, each step of the expansion *always* increases the order η , and consequently the new conditions decouple by at least one further power.* This also motivates the choice to avoid the carrier transition; it is not dependent on η , so a driving field that scales as $1/\eta$ would re-introduce the low-order terms we sought to remove.

Before moving on to find solutions to eq. (6.18), let us briefly discuss an alternative series expansion for finding the gate conditions. Instead of the frame-transformation method we have just shown, which produces a propagator as a *product* of successive terms, one could also have used the Magnus expansion to iterate *additively* towards the generator of the propagator. This is equally a valid method, but as the terms are fixed by more general Lie algebra, they are not stratified by different orders of η . In practice, their calculations tend to be more onerous, and lead to more complex functional conditions. We verified that the schemes presented in this chapter also satisfy the conditions derived from a Magnus-series approach, but the problem is more easily tractable when using the product-based series described above.

*This factor of $1/\eta$ really stumped us for a long time—we kept producing schemes that did not have the error scaling we expected. Without explicitly accounting for it, some terms that are truly of low order get mistakenly left in $\hat{\mathcal{U}}_{\text{error}}$, as they look like $f^d \eta^{d+1}$. This would be misclassified as order η^{d+1} , rather than η , so not set to zero.

6.2 Calculating solutions

The previous section gave a recipe for calculating functional conditions at each degree of η . To actually produce a gate scheme, however, we must first actually evaluate the conditions, then solve the resulting integral equations. Neither of these steps are trivial; one could calculate the conditions manually, but it is extremely tedious and error prone. Instead, this section offers a more computational approach, which I designed and implemented⁷¹.

6.2.1 Finding constraints

We begin by representing the as-yet-unknown scalar functions f by five recursive elements: a base corresponding to a specific sideband driving profile, the addition of two or more elements, the multiplication of two or more elements, the time-integral of an element, and the complex conjugation of an element. We represent a numeric prefactor as an exact fraction, and keep track of whether it is real or imaginary by a single flag. Further, we track the power of η of each term as a non-negative integer. All of the mathematical operations on scalar functions needed for eq. (6.15) can be represented in this abstract form, and used to produce concrete solutions later.

We must also track the operators that are present in the expansion. It is desirable to keep the number of unique operators tracked small, as the number of terms in each element of the sum in eq. (6.15) increases exponentially due to the multiplication. We can prune the collection by immediately discarding any term whose power of η is too high for the desired level of approximation. To go further, we need to find some *normal form* of the operators considered; we do not want to track all $\binom{n+m}{n}$ permutations of $\hat{a}^{\dagger n} \hat{a}^m$ that might arise if we can do it in fewer. We use the commutation rule

$$\hat{a}^m \hat{a}^{\dagger n} = \sum_{k=0}^{\min(n,m)} c_k \hat{a}^{\dagger n-k} \hat{a}^{m-k} \quad \text{with } c_k = \begin{cases} 1 & \text{for } k = 0 \\ \frac{(m-k+1)(n-k+1)}{k} c_{k-1} & \text{for } k > 0, \end{cases} \quad (6.21)$$

to rewrite any product of motional operators into a sum of terms in a defined order. This allows a substantial reduction; instead of the $(2n)!/(n!)^2$ different permutations of $\hat{a}^{\dagger n} \hat{a}^n$ (such as $\hat{a}^{\dagger} \hat{a} \cdots \hat{a}^{\dagger} \hat{a}$) that we might encounter, we keep track of only the n terms which have all their \hat{a}^{\dagger} operators to the left of all their \hat{a} operators.

The normal form of the qubit components of the operators is simpler. The phase convention of the driving fields was chosen in eq. (6.6) to ensure that the only qubit operator is $\hat{S}_y = \hat{\sigma}_y^{(1)} + \hat{\sigma}_y^{(2)}$. This satisfies $\hat{S}_y^3 = 4\hat{S}_y$, and so all of the qubit operators

that appear at any level of the expansion are either \hat{S}_y or \hat{S}_y^2 , with some numerical prefactor that we are already tracking. We do not need to store the operators in matrix form; a simple 3-tuple ($\{1, 2\}$, integer, integer) suffices, where the terms are the powers of \hat{S}_y , \hat{a}^\dagger , and \hat{a} respectively.

Each level of the perturbative expansion can then be stored as a hashmap of {operator: coefficient}, where the two components use the representations as described above. Taken altogether, these techniques make the two operations *find the coefficient of an operator* and *update the coefficient* constant-time complexity, while minimising the number of operators that must be tracked. This is vital for the efficiency of finding very high-order constraints, even if in practice solving the constraints from arbitrarily high orders is somewhat difficult.

With this formalism, all of the operators \hat{X} in eq. (6.18) will be of the form $\hat{a}^{\dagger p} \hat{a}^q \hat{S}_y^r$ where p and q are non-negative integers, and r is either one or two. As an illustration, consider the $\hat{a}^\dagger \hat{a} \hat{S}_y^2$ operator in the $\hat{\mathcal{U}}_2$ term with two sidebands driven. Its prefactor is*

$$i\eta^2 \text{Im}\{f_2\{f_2^*\} - 2f_1\{f_1^*\}\} = 0, \quad (6.22)$$

where the equality to zero is enforced by the requirement that there is no $\hat{a}^\dagger \hat{a} \hat{S}_y^2$ component in the Mølmer-Sørensen interaction. Physically, this term is the lowest-order error term that causes the entangling interaction to have a motion-dependent strength, even though this particular process does not change the phonon count.

6.2.2 Solving constraints

Finally, we arrive at a point where must solve a large quantity of iterated integrations to fully evaluate the constraints; recall that every appearance of the $\{f\}$ notation is an integration per eq. (6.7). One cannot reason about integral constraints with general functions, so at this point we must fix the form of the functions we will be working with. For this work, we consider only driving functions that can be parametrised as

$$f_j(t) = \sum_k c_{j,k} e^{in_{j,k}\epsilon t} \quad \text{for constant real numbers } c_{j,k} \text{ and integers } n_{j,k}, \quad (6.23)$$

with the gate occurring at a time $\tau = 2\pi/\epsilon$. Finding a solution now is equivalent to finding a set of the variables $c_{j,k}$ and $n_{j,k}$ that satisfy all the constraints. The values $n_{j,k}\epsilon$ are physically detunings from the relevant sideband, which must be kept small

*This is given in terms of the *imaginary component* operation for clarity, though programmatically we need only consider the equivalent sum and conjugation operations.

relative to the sideband separation ω_z to avoid spurious off-resonant excitation. The coefficients $c_{j,k}$ are the pulse amplitudes, related to the input laser power.

This form of the driving profiles is consistent with the standard driving of the base Mølmer–Sørensen gate, and with the multi-tone parametrisation used in chapter 5. The detunings are chosen to be integers such that the majority of constraints will be solved at the gate time simply because the integral $\int_0^{2\pi} dx e^{inx}$ is zero for non-zero integers n . Some of the constraints are not of this form, however; the entangling condition in particular is not a vanishing condition, and requires that we evaluate an integral with a zero exponent.

It is typically impractical to use general-purpose computer algebra software to calculate the vast quantity of integrals we find, as generic methods are simply too slow. We can avoid them by doing the integration more manually, though still programmatically. The group of functions with parametrisation eq. (6.23) is closed under the operations *addition*, *multiplication* and *conjugation*, but not *time integration*. One can expand the parametrisation to include a multiplying polynomial in t to keep the group closed in this case, which allows an efficient representation as

$$f_j(t) = \sum_k \sum_\ell t^\ell c_{j,k,\ell} e^{in_{j,k}\epsilon t}. \quad (6.24)$$

Non-constant polynomials are never used to drive the sidebands, only in the internal representation. The definite integrals of these functions can be calculated analytically using standard techniques, and the result is a function that is also of the form of eq. (6.24). If a scalar-function element of this group is represented by a hashmap {frequency: polynomial coefficients}, one can efficiently evaluate the necessary operations without invoking general-purpose symbolic manipulators.

At this point, we have all the ingredients to produce candidate gates. During the evaluation of the integrals, we may keep track of all the different frequencies present in integrations symbolically; they are simple summations of integers. Some, such as in terms like $\{f_1\{f_1^*\}\}$, will unavoidably become zero in the second integration, but for all others we may proceed as if they are non-zero, and build up a set of constraints on the detunings to ensure this. Finding a set of detunings—the n in eq. (6.23)—that satisfy these simple constraints is enough to solve many of the complete conditions. The amplitudes of the different tones in the driving pulses are found by a series of simultaneous equations, dependent on the particular detunings.

6.2.3 Example solutions

To improve the scaling of η beyond the base gate, one should only need to consider terms up to and including those in η^2 . Some of these terms arise from higher-order processes of the first red- and blue-sideband processes, and cannot be cancelled directly by shaping the amplitudes of these transitions. Instead, we must drive higher-order sideband transitions as well. While these more weakly couple the qubits and their motion, two factors conspire to assist us: the error processes are weak already, so we do not require much power on the higher-order sidebands; and we are concerned with the regime of increasing Lamb–Dicke parameter anyway, which improves the strength of these transitions.

Fortuitously, monochromatically driving only two pairs of sidebands is sufficient to cancel not just the factors of η^2 , but also those of η^3 . For completeness' sake, the full list of conditions that must be zero at the gate time at this order for two sidebands is:^{*}

$$\{f_1\}; \quad \{f_2\}; \quad \{\{f_1\}f_2^*\}; \quad \{f_1\{f_1\}\}; \quad \text{Re}\left\{f_1\{f_1\}\{f_2^*\} + f_1\{f_2^*\{f_1\}\}\right\}; \quad (6.25)$$

$$\{f_1\{f_2^*\}\}; \quad \{f_1\{f_2\}\}; \quad \left\{4i\{f_2\} \text{Im}\left(3f_1\{f_1^*\} - f_2\{f_2^*\}\right) + 3f_1\{f_2\{f_1^*\}\}\right\}; \quad (6.26)$$

$$\text{Re}\left\{\{f_1\}\{f_1\}f_2^*\right\}; \quad \left\{2f_1\{f_1\}\{f_1^*\} + \{f_1\}\left(f_2^*\{f_2\} - f_1^*\{f_1\}\right) - f_2\{f_2^*\{f_1\}\}\right\}; \quad (6.27)$$

$$\left\{4i\{f_2\{f_1^*\}\} \text{Im}\left(3f_1\{f_1^*\} - f_2\{f_2^*\}\right) - 2\{f_1^*\}\{f_2\}\left(3f_1^*\{f_1\} + f_2\{f_2^*\}\right) + 2f_2^*\{f_2\}\{f_2\{f_1^*\}\} - 6f_1\{f_1\}\{f_2^*\{f_2\}\} + 3f_1\{f_2\}\{f_1^*\}\{f_1^*\}\right\}; \quad (6.28)$$

$$\text{Re}\left\{6f_1\{f_1\}\{f_1^*\}\{f_2^*\{f_1\}\} - 3f_1\{f_1^*\}\{f_1^*\}\{f_2\{f_1^*\}\} + 2f_2\{f_1^*\}\{f_2^*\}\{f_2\{f_1^*\}\} - 2f_2\{f_2^*\{f_1\}\}\{f_2^*\{f_1\}\}\right\}; \quad (6.29)$$

$$\text{Im}\left\{2f_1\{f_1^*\} - f_2\{f_2^*\}\right\}. \quad (6.30)$$

The entangling condition in functional form is

$$\text{Im}\left\{f_1\{f_1^*\} + \frac{1}{2}\eta^2 f_2\{f_2^*\} - 4\eta^2\{f_1\}\left(f_2^*\{\{f_1^*\}f_2\} + f_1\{f_1^*\}^2\right)\right\} = \frac{\pi}{8}. \quad (6.31)$$

One possible solution to all of these

$$f_1(t) = \Omega \exp(2i\epsilon t) \quad \text{and} \quad f_2(t) = \Omega \exp(i\epsilon t), \quad (6.32)$$

^{*}Recall the notation $\{f\{g\}\}$ is iterated time integration $\int_0^t dt_1 f(t_1) \int_0^{t_1} dt_2 g(t_2)$.

and the gate occurring at a time $\tau = 2\pi/\epsilon$. This corresponds to a gate that makes two complete loops in phase space in order to pick up the desired entangling phase. There are multiple different possibilities for the frequencies, but this particular pair requires the least amount of total power to achieve. Regardless of the value of Ω , this satisfies all of the conditions of eqs. (6.25) to (6.29) because the two profiles are monochromatic with frequencies of 2ϵ and ϵ . The condition eq. (6.30) is satisfied because in addition to this, the amplitudes of the drivings of the two sidebands are equal. The amplitude Ω is determined by the entangling condition, which for these two profiles reduces to

$$3\eta^2 x^4 - (1 + \eta^2)x^2 + \frac{1}{8} = 0, \quad \text{where } x = \frac{\Omega}{\epsilon}. \quad (6.33)$$

We preferentially choose the smaller root to minimise power usage.

On extension to fourth order in η , the conditions become far too long to reproduce textually. They are available pre-calculated in machine-readable form in a separate repository⁷¹, along with the conditions to third order. Importantly, we were not able to find any solutions to all the conditions that use only monochromatically driven sidebands; the η^4 -dependent conditions were too intricate. Instead, we add an extra η -dependent tone onto the driving of the second sideband. With the $1/\eta$ prefactor and the natural η^2 dependence of the second sideband, this additional term contributes only processes that scale as η^2 in lowest order, and so can be used to surgically target only the problematic higher-order conditions; it does not contribute to low-order terms at all.

The lowest-power driving profiles we found that satisfy all the constraints are

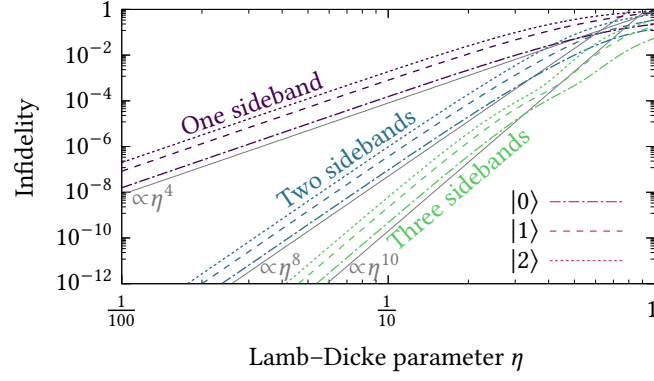
$$\begin{aligned} f_1(t) &= \Omega \exp(5i\epsilon t), & f_3(t) &= \sqrt{\frac{3}{5}} \Omega \exp(i\epsilon t), \\ f_2(t) &= \frac{\Omega}{\sqrt{5}} \left(2 \exp(2i\epsilon t) + \frac{7}{5} \frac{\Omega}{\epsilon} \eta \exp(-7i\epsilon t) \right). \end{aligned} \quad (6.34)$$

With these functions, the entangling condition to determine the ratio of $x = \Omega/\epsilon$ is

$$\frac{382}{1875} \eta^4 x^6 - \frac{56}{75} (2\eta^4 + \eta^2) x^4 + (\eta^4 + 2\eta^2 + 2) x^2 - \frac{5}{8} = 0. \quad (6.35)$$

In principle, the method we have proposed in this section allows us to go to ever higher orders. The feasibility of this in reality is limited by the available hardware control, and the difficulty in actually solving all the constraints. The computational methods presented were easily able to calculate all the constraints up to η^8 on a personal laptop within a few minutes, but actually solving them all becomes very

FIGURE 6.2—Gate infidelities of the generalised Mølmer-Sørensen scheme with one, two and three driven sidebands. For each, the infidelity is plotted for initial pure motional states with zero, one and two phonons. The scaling of the infidelity with respect to the Lamb-Dicke parameter is reduced from $\mathcal{O}(\eta^4)$ for the base gate to $\mathcal{O}(\eta^8)$ and $\mathcal{O}(\eta^{10})$ for the two- and three-sideband gates respectively. Lines proportional to these exact power laws are shown in solid grey for comparison.



difficult. We show in the next section that the scaling improvements for these first two degrees of the expansion already offer very large reductions in gate infidelities, and allow us to break out of the Lamb-Dicke regime.

6.3 Simulation results

We numerically simulate the effects of the schemes proposed by eqs. (6.32) and (6.34), to verify that the infidelity scaling with respect to η behaves as expected. For pure states of motion $|0\rangle$, $|1\rangle$ and $|2\rangle$, the gate infidelity as a function of the Lamb-Dicke parameter is shown in fig. 6.2 for the base gate, and the third- and fourth-order schemes. For these purposes, all the relevant frequencies are taken to be calibrated accurately. The simulations were done by taking the exact Hamiltonian of eq. (6.4), without any of the perturbative expansion techniques. They do not include off-resonant sideband effects, as these can be made arbitrarily small by choosing a suitably large ω_z .

The improvement in the scaling of the infidelity with respect to η is clear. All gates clearly agree with their expected scaling laws, shown in solid grey for comparison. The base gate scales approximately as η^4 , while the third- and fourth-order schemes are dominated by effects on the order of η^8 and η^{10} respectively. This is true for each of the starting pure motional states as well; one can see a non-linear dependence on the motion at the requisite order of the Lamb-Dicke parameter, but the scaling does not change. Of course, for sufficiently well excited thermal motion, one might see higher-order terms begin to dominate again. In current ion-trap labs, however, the overwhelming majority of the motional population is kept in these lower levels.

The heatmaps in fig. 6.3 show the fidelity response of the three gate schemes for highly excited thermal motional states, with mean phonon occupations up to $\bar{n} = 100$. The 99.9%-fidelity contour is presented as a guide for the range of parameters with which quantum error correction can still be achieved. The higher-

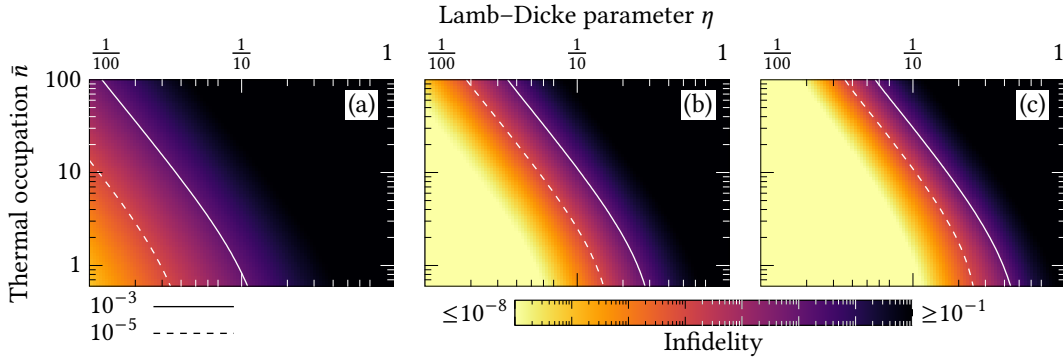


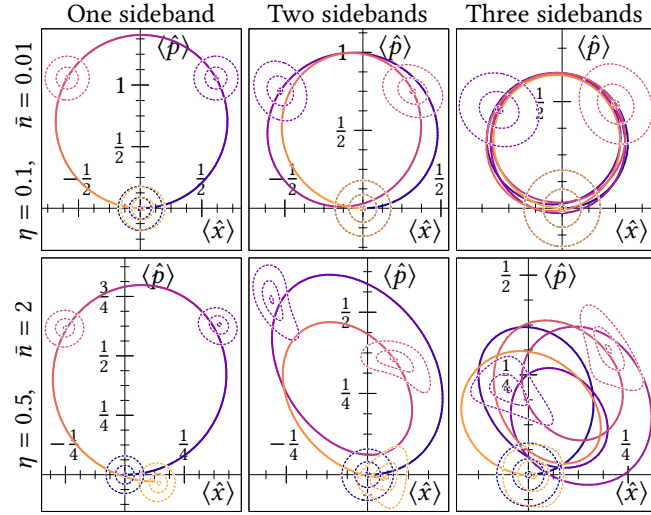
FIGURE 6.3—Heatmaps of the gate infidelity for a thermal motional state with varying mean phonon occupation and Lamb–Dicke parameter. The three plots correspond to (a) the base gate, (b) the third-order two-sideband scheme, and (c) the fourth-order three-sideband scheme. Contours denote infidelities of 10^{-3} (solid) and 10^{-5} (dashed).

order schemes can still produce valid gates even with mean thermal occupation on the order of a few phonons. This potentially allows gates to be performed that are sufficiently insensitive to the motional state that they can be performed at the Doppler-cooling limit, eschewing the need for sideband cooling.

For a quantitative comparison, prior work on producing the fastest gates in trapped ions used a Lamb–Dicke parameter $\eta \approx 1/10$, with the average motional excitation cooled to $\bar{n} \approx 1/20$ or below¹⁰⁷. The infidelity due to the breakdown of the Lamb–Dicke regime in a perfect gate with these parameters would be 1.9×10^{-4} for the base gate. The two- and three-sideband schemes as described in section 6.2.3 drastically reduce this to 1.1×10^{-7} and 7×10^{-10} respectively—several orders of magnitude for each. Of course, these fidelities are highly unlikely to be the dominant effects in any experimental realisation. Instead, we can compare the maximum achievable parameters that maintain the same error. In this case, the two-sideband scheme could maintain the same error with either a Lamb–Dicke parameter up to 0.27 or a mean thermal occupation up to 6.6, while these two numbers for the three-sideband scheme are $\eta < 0.43$ and $\bar{n} < 21$.

As in chapter 5, one can gain some insight into the behaviour of the motion throughout the gate through consideration of the phase-space displacement. This is plotted for the three schemes in question in fig. 6.4, both inside and far outside the Lamb–Dicke regime. The qubits are in the positive-value eigenstate of the \hat{S}_y operator, and the motion begins in a mixed thermal state of the given mean occupancy. The conventional driving scheme shows a substantial failure to close the phase-space loop at the gate time outside the Lamb–Dicke regime, where the new profiles have vast improvement, despite—or perhaps because of—severe deformities of their trajectories. The two- and three-sideband schemes are not single-loop gates;

FIGURE 6.4—Phase-space trajectories of the various gates (columns) both inside (top row) and far outside (bottom row) the Lamb–Dicke regime. The values of the tic marks are the same for all plots on a row, though the zooms are different. Line colour indicates time through the gate; it begins in the dark purple and ends in the light orange. The Wigner functions of the motional states at thirds of the gate time are indicated by the dashed contour lines, scaled down to 25% (top) and 5% (bottom) of their natural sizes for visibility. The squeezing effects of the higher-order sidebands can be seen via deformations of the Wigner function, but the final overlap is much better for these schemes, corresponding to reduced qubit–motion entanglement.



under perfect conditions, the number of phase-space loops is largely dictated by the ratio of the gate time to the principal detuning of the drive on the first sideband. The two-sideband profile is a two-loop gate, and the three-sideband profile is a five-loop gate. This explains the different radii of the trajectories; all three schemes enclose the same area in phase space, but the higher-order approaches repeat loops. This is not a requirement of the constraints, but in practice the solutions that are closest to constant-amplitude driving on each sideband and with the lowest power are typically multi-loop gates.

In addition to the trajectories, the Wigner functions of the motion at various points throughout the gate are indicated by contour lines. For visibility, these are scaled down around their centroids four times for the top row and twenty times for the bottom. The initial states are perfect circles, since the motion begins in a thermal state. The overlap of the initial and final Wigner states is necessary for a coherent gate. The scaling here somewhat masks the true overlap, but the gate fidelities for the bottom row of fig. 6.4 are 71%, 89% and 97.3% respectively.

Coherently applying both first-order sidebands together generates displacements in phase space. This changes to a squeezing effect for the second-order sidebands, and an asymmetric skew effect for the third-order transitions. These effects can clearly be seen in the figure; the motion for the conventional driving pattern largely retains its shape, but the new schemes produce significant distortions during the gate operation once the Lamb–Dicke parameter is large. The distortions are not seen at lower values of η , since at those coupling strengths, the higher-order processes have suppressed effects on the states. It is only when the coupling increases that these can be seen, and it is exactly these non-linear effects that are exploited to

cancel each other out at the gate time.

6.4 Additional motional modes

So far, and in our published work on this topic¹⁰⁴, we have only considered the driven motional mode. In reality, even if there are only two ions in the trap, there are more modes present. These are not directly addressed, but any transition driven on one mode implicitly drives all other modes on their carrier transition, which has its own non-linearities. We will now briefly discuss some modifications to the presented scheme that can account for these to low orders. This is introductory only; we will make several simplifying assumptions that limit the generality, as we do not yet have the tools to do this completely.

Accounting for multiple modes, labelled with subscript ℓ , the complete laser-ion interaction Hamiltonian is

$$\hat{\mathcal{H}} = f(t) \sum_j \hat{\sigma}_+^{(j)} \prod_\ell \exp\left(i\eta_{j,\ell}(\hat{a}_\ell^\dagger e^{i\kappa_\ell\omega_z t} + \hat{a}_\ell e^{-i\kappa_\ell\omega_z t})\right) + \text{H.c.} \quad (6.36)$$

The κ_ℓ are the mode-dependent modifications to the motional frequencies discussed in section 3.2 and listed in table 3.2. Using the same methods as section 6.1, this can be recast into a product of Fourier-series-like forms as

$$\hat{\mathcal{H}} = f(t) \sum_j \hat{\sigma}_+^{(j)} e^{-\frac{1}{2} \sum_\ell \eta_{j,\ell}^2} \prod_\ell \sum_{k=-\infty}^{\infty} e^{ik\kappa_\ell\omega_z t} \hat{A}_{k,\ell}(\eta_{j,\ell}) + \text{H.c.}, \quad (6.37)$$

where the $\hat{A}_{k,\ell}$ motional operators retain the same form as eq. (6.3), but act on the relevant ladder operators for the mode ℓ .

Let us assume that we are dealing with modes with incommensurate frequencies, such that there is negligible off-resonant excitation. Further, let us choose that the principal driving mode is the centre-of-mass (COM) mode, which has $\kappa_{\text{COM}} = 1$ and equal $\eta_{j,\ell} = \eta_\ell$ for all ions. If the k th transition on the COM mode is driven for qubit j , all other modes participate via their carrier transition with $k = 0$. To illustrate, if the first-order blue transition is driven, the qubit-promotion component of the applied operator, ignoring scalar prefactors, is

$$\hat{\sigma}_+^{(j)} \hat{A}_{1,\text{COM}} \prod_{\ell \neq \text{COM}} \hat{A}_{0,\ell} \approx i\hat{\sigma}_+^{(j)} \left(\eta_{\text{COM}} \hat{a}_{\text{COM}}^\dagger - \frac{1}{2} \eta_{\text{COM}}^3 \hat{a}_{\text{COM}}^{\dagger 2} \hat{a}_{\text{COM}} \right) \prod_{\ell \neq \text{COM}} \left(1 - \eta_{j,\ell}^2 \hat{a}_\ell^\dagger \hat{a}_\ell \right). \quad (6.38)$$

For modes with η_ℓ of the same order of magnitude as η_{COM} , this poses a problem: there are additional motion-dependent terms of total order η^3 that are not being

fully cancelled.

We now restrict our analysis to systems where all the relevant ions and modes have the same magnitude of $|\eta_{j,\ell}| = \eta_\ell$ for each ion within a given mode. We are dealing only with terms that are quadratic in the Lamb–Dicke parameters of the spectator modes, so the distinction between $-\eta$ and η for different ions is irrelevant. Because the spectator carrier transitions only introduce even powers of η , we seek to cancel out these processes by application of the second-order sidebands on the spectator modes. We wish to avoid the first-order sidebands in this case because their inclusion will cause various *odd* powers of η to appear, which would need to be handled individually. With the second-order sidebands, any higher-order processes from the driven transitions will quickly rise to orders of η that we are already neglecting.

The perturbative expansion method of section 6.1 can now be applied, with minor modifications. When finding the lowest-order terms in η , we now consider the order of a term $\eta_1^a \eta_2^b$ to be $a + b$, or with some alternate weighting for the different modes, if desired. We must also modify our prefactors for the different driving profiles of the sidebands. The new form of eq. (6.5) is

$$f(t) = -ie^{\sum_\ell \eta_\ell^2/2} \sum_\ell \frac{1}{\eta_\ell} \sum_{k>0} \left(f_{\ell,k}(t) e^{-ik\kappa_\ell \omega_z t} + (-1)^k f_{\ell,k}^*(t) e^{ik\kappa_\ell \omega_z t} \right). \quad (6.39)$$

We include the factors of $1/\eta_\ell$ separately, and applied only to transitions on the relevant mode, though the total exponential prefactor is always applicable.

Without any correction, when considering terms up to a total order of η^3 , each spectator mode contributes a thermal dependence via the operator $\hat{S}_y \eta_\ell^2 \hat{a}_\ell^\dagger \hat{a}_\ell$. These terms can be cancelled by driving the relevant modes on their second sidebands. This would introduce further thermal cross-terms from higher-order processes, but at the level of η^3 , only the leading-order process contributes, which we can use to cancel the unwanted spectator carrier effects.

Using the same programmatic method as in the previous sections, but updated to symbolically track the multiple modes, the conditions for nulling non-linear effects were found to low order. These can be solved by the driving profiles

$$f_{\text{COM},1}(t) = \Omega \exp(2i\epsilon t), \quad f_{\text{COM},2}(t) = \Omega \exp(i\epsilon t), \quad \text{and} \quad f_{\ell,2}(t) = \Omega \exp(i\epsilon t). \quad (6.40)$$

The updated entangling condition for these requires that $x = \Omega/\epsilon$ is a root of

$$3\eta_{\text{COM}}^2 x^4 - \left(1 + \sum_\ell \eta_\ell^2\right) x^2 + \frac{1}{8} = 0, \quad (6.41)$$

where the summation is over all driven modes, including the COM mode. Note that despite immediate appearances, this scheme does not drive the second sidebands of all modes with the same power; each is scaled by the factor of $1/\eta_\ell$ in eq. (6.39). Even the power of the transitions on the COM mode are slightly modified from the earlier scheme, due to the different prefactors between eqs. (6.5) and (6.39). This latter difference is relatively slight, however.

The methods of this chapter should, in principle, allow us to cancel out higher-order effects of the carrier transitions from spectator modes, with orders η^4 and beyond. While I was able to calculate the functional conditions for an ideal gate at this level of approximation, actually solving all of them simultaneously is rather difficult. There is no particular reason to assume that it is impossible, but more insight into the potential forms of the solutions might be needed to make further progress.

6.5 Outlook

This chapter has developed a system for producing trapped-ion gates that go beyond the non-linear, weakly coupled regime. This has been a fundamental limit on the achievable fidelities and speeds in state-of-the-art trapped-ion quantum gates, as it could never be overcome for conventional gates, no matter the quality or precision of the experimental control apparatus. The method described here was a systematic approach to cancelling contributions from undesired non-linear processes, removing them order-by-order in the coupling strength. In principle, the methods described here can produce extremely high fidelity two-qubit gates even without protracted sideband cooling of the various motional modes. This opens avenues to both high-temperature and high-speed quantum information processing with trapped ions, with less reset time between different shots of circuits. Cooling cycles and reset times are among the longest operations in practical ion-trap quantum computing, and reducing their necessity is an important contribution, let alone the possible improvements in fidelity scaling.

The method produces a set of functional constraints, and does not dictate the exact form of the solutions. Simulations in this chapter focussed on the simplest, lowest-power solutions, where each sideband is driven as close to monochromatically as possible. This is easily achievable for modern control software, but is not a true necessity. Mathematically, it is simple to combine this gate design with many of the current pulse-shaping techniques to make gates robust against drifting calibrations. This includes the multi-tone gates described in chapter 5 to decouple the fidelity

from qubit-frequency drifts, but also the motional-drift-resilient gates achieved in other works by amplitude shaping^{42,110,132} or pulse shaping⁸⁰.

Let us take the multi-tone schemes of the previous chapter as an example. One would need to make minor modifications to the numerical optimisations to ensure that all the incommensurate-frequency conditions from eqs. (6.25) to (6.29) are maintained, but this can be done manually, simply by choosing the detunings used on the tones. The entangling condition would already have been satisfied by the optimiser, and all that is left is to fix the final condition of eq. (6.30). This can be done—in theory—with a monochromatic driving profile on the second sideband; it is relatively trivial to carry out the integrations and simply set the amplitude to the correct value.

It should be noted, however, that the driving on the higher-order sidebands may not itself be decoupled from the same miscalibrations if done in this manner. This would be more noticeable in the analytically-derived multitone schemes to make gates resilient against *motional* drift^{110,132}. In these, one derives the exact amplitudes of the different tones by expanding the time-evolution operator in terms of the frequency error, and using the method of Lagrange multipliers to minimise the error components. This approach is not as straightforward when multiple sidebands are involved, but is feasible in principle. One could use the methods given in this chapter to approximate the generators of the propagators, and then expand each as a Taylor series to produce an—exceedingly complex—expression for the final fidelity. This could then be approached with the same techniques as in previous work to find a complete scheme that is both non-linear and robust against miscalibrations of the motional mode.

Further, while we have discussed methods to remove the non-linearities entering from unwanted participation of spectator modes, so far these are limited to gates performed on ions that participate equally in all the considered modes. The normal modes in table 3.2 show that few pairs of ions have the same coupling to all possible modes once there are more than two ions in a trap. It is an open question whether this scheme could be extended to account for, say, the breathing mode of two neighbouring ions on one side of a four-ion trap. It is not immediately clear how to handle the unequal participation. In eq. (6.37), we still assumed a global irradiation field for the ions. One potential avenue of research is to address the ions individually, with suitably different laser amplitudes—one can imagine a more complex set of conditions where even the unequal mode participation might be overcome. As interesting as it is, however, further work in this vein is beyond the scope of this thesis.

Chapter 7

Conclusion

My PhD research has been an investigation into the creation and verification of robust coherence and entanglement in trapped ions. The techniques presented here have all been working towards a goal of greater, more resilient coherent control of the quantum behaviour of trapped ions, with a view to enhancing noisy intermediate-scale quantum information processing in this medium. Going further, I hope that the work here is one more (small) stepping stone along the path to large-scale quantum computation, and the realisation of new models of computing.

My work in chapter 4 demonstrated that a general certifier for multilevel coherence could be made robust against false positives, even when the basis of coherence was not directly accessible to measurement. This went beyond previous work on the subject, allowing even untrusted coherent manipulations to be made by mapping a set of basis states onto a smaller, physically distinct Hilbert space, without risking the certifier incorrectly claiming the presence of high-order coherence. We then demonstrated this protocol by numerically finding mapping sequences that could generalise the Ramsey coherence experiment to higher orders, coherently mapping two orthogonal subspaces of motional basis states onto different qubit states for measurement. With the experimental team at Imperial, we implemented this in real hardware, successfully certifying 3-coherence in the motional state of a trapped ion, and offering strong evidence in favour of having created 4-coherence. This method required only a one-dimensional interference pattern, significantly reducing the number of measurements and experimental complexity necessary to certify coherence.

Moving to entanglement, chapter 5 presented a modification to the Mølmer-Sørensen scheme to make two-qubit entangling operations in trapped ions more resilient to fluctuations in the frequencies of the qubits. This used a multi-tone approach that has previously found success at decoupling the gate fidelities from errors in the motional frequency. We saw order-of-magnitude improvements in the infidelities of the two-qubit gate at errors that would usually cause a system to leave the fault-tolerance region. This opens up alternative avenues for the

managing errors in real experimental situations, for example in larger-scale systems of magnetic-field-sensitive qubits where it becomes near-impossible to completely control all fluctuating fields over the trapping zones. Still, though, the improvements seen here were perhaps not as marked as when the same multi-tone parametrisation is applied to different classes of error^{110,132}, so it seems possible that a more fruitful strategy here would be to use this scheme to become robust against motional errors, and use some other decoupling method to protect against qubit miscalibrations, such as single-qubit dynamic decoupling⁷⁶.

The most exciting result for the future of ion traps is the method in chapter 6 to bring trapped-ion quantum computing outside the weakly coupled linear regime it has been confined to since its inception. We overcame a previously fundamental restriction on the infidelities of trapped-ion entangling gates, and demonstrated a systematic method for decoupling the gate operation from the motion to ever-higher orders. We found in exact simulations that our scheme brought about improvements in the asymptotic power-law scaling of the gate infidelity with respect to the ion–motion coupling. Using the trap parameters of the current state of the art for ion-trap entangling gates¹⁰⁷, we showed that our scheme can in theory support a 2000-fold improvement in the linear-regime error simply by driving a single additional pair of sidebands, and this increases by several more orders of magnitude if another pair of transitions are included.

This new multi-sideband scheme opens many new areas of research. Without modification, it paves the way to trapped-ion quantum computing with hot motion; state-of-the-art fidelities can be reached, even if the ions are not cooled beyond the Doppler limit. Further, the method is a systematic way to derive functional constraints on the driving profiles of different sidebands, which should allow it to be used in conjunction with existing pulse-shaping methods of making robust gates, including the multi-tone methods discussed in chapter 5. We sketched the procedure for unifying a resilient gate scheme with this new, non-linear approach to entanglement generation.

Additional theoretical work in this vein may also be in cancelling out the non-linearities that enter from spectator motional modes. We derived a set of constraints that decouple all relevant modes in which the ions participate equally, such that the resulting infidelity scaling is of order η^8 . We sketched out possible paths forward to extend this to higher orders of the coupling strength, or to account for motional modes with unequal ion participation. We also look forwards to the results of an ongoing experimental collaboration with the ion-trapping group at Imperial, who hope to implement our schemes in the coming months.

Ultimately, the true proof of all of these methods is in experimental realisations. We do not yet know for certain if a quantum advantage is possible in computational tasks, nor if trapped ions will be the best platform for larger-scale quantum computation. Still, all of the results presented here are improvements in current techniques for characterisation and generation of quantum behaviour in ion traps. There is a long way to go if we are to achieve general-purpose quantum computation, and the work here represents one more step in that journey.

Bibliography

- [1] J. Åberg, *Quantifying Superposition*, [arXiv:quant-ph/0612146](#) (2006). Cited on page 43.
- [2] I. Arrazola, J. Casanova, J. S. Pedernales, Z.-Y. Wang, E. Solano, and M. B. Plenio, *Pulsed dynamical decoupling for fast and robust two-qubit gates on trapped ions*, [Physical Review A](#) **97**, 052312 (2018). Cited on page 74.
- [3] F. Arute *et al.*, *Quantum supremacy using a programmable superconducting processor*, [Nature](#) **574**, 505 (2019). Cited on page 11.
- [4] 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), [arXiv:1512.04600](#). Cited on pages 10, 23, 38, 74, and 91.
- [5] P. A. Barton, C. J. S. Donald, D. M. Lucas, D. A. Stevens, A. M. Steane, and D. N. Stacey, *Measurement of the Lifetime of the $3d^2D_{\frac{5}{2}}$ State in $^{40}\text{Ca}^+$* , [Physical Review A](#) **62**, 032503 (2000). Cited on page 25.
- [6] T. Baumgratz, M. Cramer, and M. B. Plenio, *Quantifying Coherence*, [Physical Review Letters](#) **113**, 140401 (2014), [arXiv:1311.0275](#). Cited on pages 41 and 43.
- [7] A. Ben-Kish, B. DeMarco, V. Meyer, M. Rowe, J. Britton, W. M. Itano, B. M. Jelenković, C. Langer, D. Leibfried, T. Rosenband, and D. J. Wineland, *Experimental Demonstration of a Technique to Generate Arbitrary Quantum Superposition States of a Harmonically Bound Spin-1/2 Particle*, [Physical Review Letters](#) **90**, 037902 (2003). Cited on page 53.
- [8] J. Benhelm, G. Kirchmair, C. F. Roos, and R. Blatt, *Towards fault-tolerant quantum computing with trapped ions*, [Nature Physics](#) **4**, 463 (2008), [arXiv:0803.2798](#). Cited on pages 26, 74, and 75.
- [9] 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 (1980). Cited on page 10.
- [10] P. Benioff, *Quantum mechanical Hamiltonian models of Turing machines*, [Journal of Statistical Physics](#) **29**, 515 (1982). Cited on page 10.
- [11] C. H. Bennett, G. Brassard, S. Popescu, B. Schumacher, J. A. Smolin, and W. K. Wootters, *Purification of Noisy Entanglement and Faithful Teleportation via Noisy Channels*, [Physical Review Letters](#) **76**, 722 (1996). Cited on page 41.
- [12] A. Bermudez, P. O. Schmidt, M. B. Plenio, and A. Retzker, *Robust trapped-ion quantum logic gates by continuous dynamical decoupling*, [Physical Review A](#) **85**, 040302(R) (2012), [arXiv:1110.1870](#). Cited on page 75.
- [13] A. Bermudez, X. Xu, R. Nigmatullin, J. O’Gorman, V. Negnevitsky, P. Schindler, T. Monz, U. G. Poschinger, C. Hempel, J. Home, F. Schmidt-Kaler, M. Biercuk, R. Blatt, S. Benjamin, and M. Müller, *Assessing the Progress of Trapped-Ion Processors Towards Fault-Tolerant Quantum Computation*, [Physical Review X](#) **7**, 041061 (2017), [arXiv:1705.02771](#). Cited on page 23.
- [14] R. Blümel, N. Grzesiak, N. H. Nguyen, A. M. Green, M. Li, A. Maksymov, N. M. Linke, and Y. Nam, *Efficient Stabilized Two-Qubit Gates on a Trapped-Ion Quantum Computer*, [Physical Review Letters](#) **126**, 220503 (2021). Cited on pages 11 and 89.
- [15] D. Bluvstein, H. Levine, G. Semeghini, T. T. Wang, S. Ebadi, M. Kalinowski, A. Keesling, N. Maskara, H. Pichler, M. Greiner, V. Vuletić, and M. D. Lukin, *A quantum processor based on coherent transport of entangled atom arrays*, [Nature](#) **604**, 451 (2022). Cited on page 11.
- [16] A. Bonfiglioli and R. Fulci, *Topics in Noncommutative Algebra: The Theorem of Campbell, Baker, Hausdorff and Dynkin*, Lecture Notes in Mathematics, Vol. 2034 (Springer-Verlag, Berlin, Heidelberg, 2012). Cited on page 20.

Bibliography

- [17] B. H. Bransden and C. J. Joachain, *Physics of Atoms and Molecules* (Longman, London ; New York, 1983). Cited on page 30.
- [18] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, Oxford ; New York, 2002). Cited on page 21.
- [19] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage, *Trapped-ion quantum computing: Progress and challenges*, *Applied Physics Reviews* **6**, 021314 (2019), [arXiv:1904.04178](#). Cited on page 75.
- [20] A. Castellini, R. Lo Franco, L. Lami, A. Winter, G. Adesso, and G. Compagno, *Indistinguishability-enabled coherence for quantum metrology*, *Physical Review A* **100**, 012308 (2019), [arXiv:1903.10582](#). Cited on page 42.
- [21] L. H. Chan, K. Chen, C. Li, C. W. Wong, and C. Y. Yau, *On higher-order moment and cumulant estimation*, *Journal of Statistical Computation and Simulation* **90**, 747 (2020). Cited on page 61.
- [22] E. Chitambar and G. Gour, *Quantum resource theories*, *Reviews of Modern Physics* **91**, 025001 (2019). Cited on pages 16 and 43.
- [23] T. Choi, S. Debnath, T. A. Manning, C. Figgatt, Z.-X. Gong, L.-M. Duan, and C. Monroe, *Optimal Quantum Control of Multimode Couplings between Trapped Ion Qubits for Scalable Entanglement*, *Physical Review Letters* **112**, 190502 (2014), [arXiv:1401.1575](#). Cited on page 74.
- [24] I. L. Chuang, N. Gershenfeld, and M. Kubinec, *Experimental Implementation of Fast Quantum Searching*, *Physical Review Letters* **80**, 3408 (1998). Cited on page 10.
- [25] J. I. Cirac and P. Zoller, *Quantum Computations with Cold Trapped Ions*, *Physical Review Letters* **74**, 4091 (1995), [arXiv:quant-ph/0305129](#). Cited on pages 36 and 74.
- [26] R. Cools, *An encyclopaedia of cubature formulas*, *Journal of Complexity* **19**, 445 (2003). Cited on page 84.
- [27] R. Cools and P. Rabinowitz, *Monomial cubature rules since “Stroud”: A compilation*, *Journal of Computational and Applied Mathematics* **48**, 309 (1993). Cited on page 84.
- [28] O. Corfield, *Quantum Coherence in Trapped Ions*, Ph.D. thesis, Imperial College London (2022). Cited on pages 27, 32, 64, and 68.
- [29] O. Corfield, J. Lishman, C. Lee, J. Mosca Toba, G. Porter, J. M. Heinrich, S. C. Webster, F. Mintert, and R. C. Thompson, *Certifying Multilevel Coherence in the Motional State of a Trapped Ion*, *PRX Quantum* **2**, 040359 (2021), [arXiv:2106.12939](#). Cited on pages 2, 12, and 41.
- [30] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, and C. Monroe, *Demonstration of a small programmable quantum computer with atomic qubits*, *Nature* **536**, 63 (2016), [arXiv:1603.04512](#). Cited on pages 25 and 74.
- [31] H. Dehmelt, *Radiofrequency Spectroscopy of Stored Ions I: Storage*, in *Advances in Atomic and Molecular Physics*, Vol. 3 (Elsevier, 1968) pp. 53–72. Cited on page 23.
- [32] D. Deutsch, *Quantum Theory, the Church-Turing Principle and the Universal Quantum Computer*, *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences* **400**, 97 (1985). Cited on page 10.
- [33] D. Deutsch, *Quantum Computational Networks*, *Proceedings of the Royal Society A* **425**, 73 (1989). Cited on page 10.
- [34] B. Dive, N. Koukoulekidis, S. Mousafeiris, and F. Mintert, *Characterisation of multi-level quantum coherence without ideal measurements*, *Physical Review Research* **2**, 013220 (2020), [arXiv:1901.08599](#). Cited on pages 42, 45, 46, 47, 50, 51, and 71.
- [35] D. P. DiVincenzo, *The Physical Implementation of Quantum Computation*, *Fortschritte der Physik* **48**, 771 (2000), [arXiv:quant-ph/0002077](#). Cited on pages 11 and 24.
- [36] D. P. DiVincenzo and P. W. Shor, *Fault-Tolerant Error Correction with Efficient Quantum Codes*, *Physical Review Letters* **77**, 3260 (1996). Cited on page 11.
- [37] W. Dür and H. J. Briegel, *Entanglement purification and quantum error correction*, *Reports on Progress in Physics* **70**, 1381 (2007). Cited on page 16.

Bibliography

- [38] J. Eschner, G. Morigi, F. Schmidt-Kaler, and R. Blatt, *Laser cooling of trapped ions*, *Journal of the Optical Society of America B* **20**, 1003 (2003). Cited on page 24.
- [39] R. P. Feynman, *Simulating physics with computers*, *International Journal of Theoretical Physics* **21**, 467 (1982). Cited on page 10.
- [40] 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), [arXiv:1604.00032](#). Cited on pages 10, 23, 38, 74, and 91.
- [41] S. A. Gardiner, J. I. Cirac, and P. Zoller, *Nonclassical states and measurement of general motional observables of a trapped ion*, *Physical Review A* **55**, 1683 (1997). Cited on page 53.
- [42] F. Haddadfarshi and F. Mintert, *High fidelity quantum gates of trapped ions in the presence of motional heating*, *New Journal of Physics* **18**, 123007 (2016). Cited on pages 75, 79, 88, and 109.
- [43] H. Häffner, S. Gulde, M. Riebe, G. Lancaster, C. Becher, J. Eschner, F. Schmidt-Kaler, and R. Blatt, *Precision Measurement and Compensation of Optical Stark Shifts for an Ion-Trap Quantum Processor*, *Physical Review Letters* **90**, 4 (2003), [arXiv:physics/0212040](#). Cited on page 31.
- [44] E. L. Hahn, *Spin Echoes*, *Physical Review* **80**, 580 (1950). Cited on page 74.
- [45] P. C. Haljan, P. J. Lee, K.-A. Brickman, M. Acton, L. Deslauriers, and C. Monroe, *Entanglement of trapped-ion clock states*, *Physical Review A* **72**, 062316 (2005). Cited on page 74.
- [46] B. C. Hall, *Lie Groups, Lie Algebras, and Representations: An Elementary Introduction*, 2nd ed., Graduate Texts in Mathematics, Vol. 222 (Springer International Publishing, Cham, 2015). Cited on pages 20 and 95.
- [47] C. R. Harris *et al.*, *Array programming with NumPy*, *Nature* **585**, 357 (2020). Cited on page 12.
- [48] 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), [arXiv:1403.1524](#). Cited on pages 25 and 74.
- [49] 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), [arXiv:1606.08409](#). Cited on pages 25, 26, and 75.
- [50] M. Harvey, *Imperial College Research Computing Service* (2017). Cited on pages 12 and 51.
- [51] M. Hillery, *Coherence as a resource in decision problems: The Deutsch-Jozsa algorithm and a variation*, *Physical Review A* **93**, 012111 (2016). Cited on page 41.
- [52] J. P. Home, M. J. McDonnell, D. M. Lucas, G. Imreh, B. C. Keitch, D. J. Szwer, N. R. Thomas, S. C. Webster, D. N. Stacey, and A. M. Steane, *Deterministic entanglement and tomography of ion-spin qubits*, *New Journal of Physics* **8**, 188 (2006). Cited on page 25.
- [53] C. K. Hong, Z. Y. Ou, and L. Mandel, *Measurement of subpicosecond time intervals between two photons by interference*, *Physical Review Letters* **59**, 2044 (1987). Cited on pages 14 and 42.
- [54] M. Horodecki, P. Horodecki, and R. Horodecki, *Separability of mixed states: Necessary and sufficient conditions*, *Physics Letters A* **223**, 1 (1996). Cited on page 41.
- [55] R. Horodecki, *Quantum Information*, *Acta Physica Polonica A* **139**, 197 (2021). Cited on page 41.
- [56] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, *Quantum entanglement*, *Reviews of Modern Physics* **81**, 865 (2009). Cited on pages 16 and 41.
- [57] P. Hrmo, *Ground State Cooling of the Radial Motion of a Single Ion in a Penning Trap and Coherent Manipulation of Small Numbers of Ions*, *Ph.D. thesis*, Imperial College London (2018). Cited on pages 25 and 27.

Bibliography

- [58] A. C. Hughes, V. M. Schäfer, K. Thirumalai, D. P. Nadlinger, S. R. Woodrow, D. M. Lucas, and C. J. Ballance, *Benchmarking a High-Fidelity Mixed-Species Entangling Gate*, [*Physical Review Letters* **125**, 080504 \(2020\)](#). Cited on page 23.
- [59] D. F. V. James, *Quantum dynamics of cold trapped ions with application to quantum computation*, [*Applied Physics B: Lasers and Optics* **66**, 181 \(1998\)](#), [arXiv:quant-ph/9702053](#). Cited on page 32.
- [60] V. Jarlaud, *Sideband Cooling of Ion Coulomb Crystals in a Penning Trap*, [Ph.D. thesis](#), Imperial College London (2018). Cited on page 25.
- [61] J. Johansson, P. Nation, and F. Nori, *QuTiP 2: A Python framework for the dynamics of open quantum systems*, [*Computer Physics Communications* **184**, 1234 \(2013\)](#), [arXiv:1211.6518](#). Cited on pages 12 and 79.
- [62] J. A. Jones, M. Mosca, and R. H. Hansen, *Implementation of a quantum search algorithm on a quantum computer*, [*Nature* **393**, 344 \(1998\)](#). Cited on page 10.
- [63] M. Joshi, *Coherent Dynamics of Trapped Ions Within and Outside the Lamb-Dicke Regime*, [Ph.D. thesis](#), Imperial College London (2018). Cited on page 25.
- [64] E. Knill, *Quantum computing with realistically noisy devices*, [*Nature* **434**, 39 \(2005\)](#), [arXiv:quant-ph/0410199](#). Cited on pages 79 and 86.
- [65] P. Kok, W. J. Munro, K. Nemoto, T. C. Ralph, J. P. Dowling, and G. J. Milburn, *Linear optical quantum computing with photonic qubits*, [*Reviews of Modern Physics* **79**, 135 \(2007\)](#), [arXiv:quant-ph/0512071](#). Cited on page 11.
- [66] K. Korzekwa, M. Lostaglio, J. Oppenheim, and D. Jennings, *The extraction of work from quantum coherence*, [*New Journal of Physics* **18**, 023045 \(2016\)](#), [arXiv:1506.07875](#). Cited on page 41.
- [67] B. P. Lanyon, T. J. Weinhold, N. K. Langford, M. Barbieri, D. F. V. James, A. Gilchrist, and A. G. White, *Experimental Demonstration of a Compiled Version of Shor's Algorithm with Quantum Entanglement*, [*Physical Review Letters* **99**, 250505 \(2007\)](#). Cited on page 10.
- [68] 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**, e1601540 \(2017\)](#), [arXiv:1508.00420](#). Cited on page 75.
- [69] F. Levi and F. Mintert, *A quantitative theory of coherent delocalization*, [*New Journal of Physics* **16**, 033007 \(2014\)](#). Cited on pages 41 and 43.
- [70] J. Lishman, [Code accompaniment to \(Corfield, Lishman et al., 2021\)](#) (2021). Cited on page 58.
- [71] J. Lishman, [Code accompaniment to \(Sameti, Lishman and Mintert, 2021\)](#) (2021). Cited on pages 98 and 102.
- [72] J. Lishman and F. Mintert, *Trapped-ion entangling gates robust against qubit frequency errors*, [*Physical Review Research* **2**, 033117 \(2020\)](#), [arXiv:2004.07253](#). Cited on pages 2, 12, and 74.
- [73] R. Loudon, *The Quantum Theory of Light*, 3rd ed., Oxford Science Publications (Oxford University Press, Oxford ; New York, 2000). Cited on page 30.
- [74] C.-Y. Lu, D. E. Browne, T. Yang, and J.-W. Pan, *Demonstration of a Compiled Version of Shor's Quantum Factoring Algorithm Using Photonic Qubits*, [*Physical Review Letters* **99**, 250504 \(2007\)](#). Cited on page 10.
- [75] W. Magnus, *On the exponential solution of differential equations for a linear operator*, [*Communications on Pure and Applied Mathematics* **7**, 649 \(1954\)](#). Cited on page 20.
- [76] T. Manovitz, A. Rotem, R. Shaniv, I. Cohen, Y. Shapira, N. Akerman, A. Retzker, and R. Ozeri, *Fast Dynamical Decoupling of the Mølmer-Sørensen Entangling Gate*, [*Physical Review Letters* **119**, 220505 \(2017\)](#). Cited on pages 89 and 111.
- [77] T. Manovitz, Y. Shapira, L. Gazit, N. Akerman, and R. Ozeri, *Trapped-Ion Quantum Computer with Robust Entangling Gates and Quantum Coherent Feedback*, [*PRX Quantum* **3**, 010347 \(2022\)](#). Cited on page 90.

Bibliography

- [78] D. Manzano, *A short introduction to the Lindblad master equation*, *AIP Advances* **10**, 025106 (2020). Cited on page 21.
- [79] G. J. Milburn, S. Schneider, and D. F. V. James, *Ion Trap Quantum Computing With Warm Ions*, *Fortschritte der Physik* **48**, 801 (2000). Cited on page 88.
- [80] A. R. Milne, C. L. Edmunds, C. Hempel, F. Roy, S. Mavadia, and M. J. Biercuk, *Phase-Modulated Entangling Gates Robust to Static and Time-Varying Errors*, *Physical Review Applied* **13**, 024022 (2020), [arXiv:1808.10462](#). Cited on pages 75, 89, and 109.
- [81] F. Mintert and C. Wunderlich, *Ion-Trap Quantum Logic Using Long-Wavelength Radiation*, *Physical Review Letters* **87**, 257904 (2001), [arXiv:quant-ph/0104041](#). Cited on page 25.
- [82] K. Mølmer and A. Sørensen, *Multiparticle Entanglement of Hot Trapped Ions*, *Physical Review Letters* **82**, 1835 (1999). Cited on page 38.
- [83] C. Monroe, D. M. Meekhof, B. E. King, W. M. Itano, and D. J. Wineland, *Demonstration of a Fundamental Quantum Logic Gate*, *Physical Review Letters* **75**, 4714 (1995). Cited on pages 10, 25, and 36.
- [84] C. Monroe, R. Raussendorf, A. Ruthven, K. R. Brown, P. Maunz, L.-M. Duan, and J. Kim, *Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects*, *Physical Review A* **89**, 022317 (2014). Cited on page 75.
- [85] T. Monz, P. Schindler, J. T. Barreiro, M. Chwalla, D. Nigg, W. A. Coish, M. Harlander, W. Hänsel, M. Hennrich, and R. Blatt, *14-Qubit Entanglement: Creation and Coherence*, *Physical Review Letters* **106**, 130506 (2011). Cited on page 26.
- [86] A. H. Myerson, D. J. Szwer, S. C. Webster, D. T. C. Allcock, M. J. Curtis, G. Imreh, J. A. Sherman, D. N. Stacey, A. M. Steane, and D. M. Lucas, *High-Fidelity Readout of Trapped-Ion Qubits*, *Physical Review Letters* **100**, 200502 (2008). Cited on page 25.
- [87] W. Nagourney, J. Sandberg, and H. Dehmelt, *Shelved optical electron amplifier: Observation of quantum jumps*, *Physical Review Letters* **56**, 2797 (1986). Cited on page 25.
- [88] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th ed. (Cambridge University Press, Cambridge ; New York, 2010). Cited on page 14.
- [89] D. K. L. Oi and J. Åberg, *Fidelity and Coherence Measures from Interference*, *Physical Review Letters* **97**, 220404 (2006), [arXiv:quant-ph/0603157](#). Cited on page 44.
- [90] S. Olmschenk, K. C. Younge, D. L. Moehring, D. N. Matsukevich, P. Maunz, and C. Monroe, *Manipulation and detection of a trapped Yb^+ hyperfine qubit*, *Physical Review A* **76**, 052314 (2007). Cited on page 25.
- [91] W. Paul, *Electromagnetic traps for charged and neutral particles*, *Reviews of Modern Physics* **62**, 531 (1990). Cited on page 23.
- [92] A. Peres, *Separability Criterion for Density Matrices*, *Physical Review Letters* **77**, 1413 (1996). Cited on pages 41 and 42.
- [93] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, M. S. Allman, C. H. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, C. Ryan-Anderson, and B. Neyenhuis, *Demonstration of the trapped-ion quantum CCD computer architecture*, *Nature* **592**, 209 (2021). Cited on pages 11, 25, and 89.
- [94] I. Pogorelov *et al.*, *Compact Ion-Trap Quantum Computing Demonstrator*, *PRX Quantum* **2**, 020343 (2021). Cited on page 26.
- [95] U. G. Poschinger, G. Huber, F. Ziesel, M. Deiß, M. Hettrich, S. A. Schulz, K. Singer, G. Poulsen, M. Drewsen, R. J. Hendricks, and F. Schmidt-Kaler, *Coherent manipulation of a $^{40}\text{Ca}^+$ spin qubit in a micro ion trap*, *Journal of Physics B: Atomic, Molecular and Optical Physics* **42**, 154013 (2009). Cited on page 25.
- [96] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes*, 3rd ed. (Cambridge University Press, 2007). Cited on pages 50, 79, 80, 83, and 84.

Bibliography

- [97] S. Rana, P. Parashar, and M. Lewenstein, *Trace-distance measure of coherence*, [Physical Review A **93**, 012110 \(2016\)](#). Cited on page 43.
- [98] M. Riebe, K. Kim, P. Schindler, T. Monz, P. O. Schmidt, T. K. Körber, W. Hänsel, H. Häffner, C. F. Roos, and R. Blatt, *Process Tomography of Ion Trap Quantum Gates*, [Physical Review Letters **97**, 220407 \(2006\)](#), [arXiv:quant-ph/0609228](#). Cited on page 26.
- [99] M. Ringbauer, T. R. Bromley, M. Cianciaruso, L. Lami, W. Y. S. Lau, G. Adesso, A. G. White, A. Fedrizzi, and M. Piani, *Certification and Quantification of Multilevel Quantum Coherence*, [Physical Review X **8**, 041007 \(2018\)](#), [arXiv:1707.05282](#). Cited on pages 42 and 43.
- [100] J. Roffe, *Quantum error correction: An introductory guide*, [Contemporary Physics **60**, 226 \(2019\)](#). Cited on page 11.
- [101] C. F. Roos, *Ion trap quantum gates with amplitude-modulated laser beams*, [New Journal of Physics **10**, 013002 \(2008\)](#), [arXiv:0710.1204](#). Cited on pages 37 and 92.
- [102] 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**, 254 \(2016\)](#). Cited on page 25.
- [103] C. A. Sackett, D. Kielpinski, B. E. King, C. Langer, V. Meyer, C. J. Myatt, M. Rowe, Q. A. Turchette, W. M. Itano, D. J. Wineland, and C. Monroe, *Experimental entanglement of four particles*, [Nature **404**, 256 \(2000\)](#). Cited on pages 38 and 75.
- [104] M. Sameti, J. Lishman, and F. Mintert, *Strong-coupling quantum logic of trapped ions*, [Physical Review A **103**, 052603 \(2021\)](#), [arXiv:2003.11718](#). Cited on pages 2, 12, 91, and 106.
- [105] J. P. Santos, L. C. Céleri, G. T. Landi, and M. Paternostro, *The role of quantum coherence in non-equilibrium entropy production*, [npj Quantum Information **5**, 23 \(2019\)](#), [arXiv:1707.08946](#). Cited on page 41.
- [106] T. Sauter, W. Neuhauser, R. Blatt, and P. E. Toschek, *Observation of Quantum Jumps*, [Physical Review Letters **57**, 1696 \(1986\)](#). Cited on page 25.
- [107] V. M. Schäfer, C. J. Ballance, K. Thirumalai, L. J. Stephenson, T. G. Ballance, A. M. Steane, and D. M. Lucas, *Fast quantum logic gates with trapped-ion qubits*, [Nature **555**, 75 \(2018\)](#), [arXiv:1709.06952](#). Cited on pages 23, 26, 32, 75, 91, 104, and 111.
- [108] F. Schmidt-Kaler, H. Häffner, S. Gulde, M. Riebe, G. P. T. Lancaster, T. Deuschle, C. Becher, W. Hänsel, J. Eschner, C. F. Roos, and R. Blatt, *How to realize a universal quantum gate with trapped ions*, [Applied Physics B **77**, 789 \(2003\)](#), [arXiv:quant-ph/0312162](#). Cited on page 25.
- [109] 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 \(2003\)](#). Cited on pages 25 and 36.
- [110] Y. Shapira, R. Shaniv, T. Manovitz, N. Akerman, and R. Ozeri, *Robust entanglement gates for trapped-ion qubits*, [Physical Review Letters **121**, 180502 \(2018\)](#), [arXiv:1805.06806](#). Cited on pages 75, 79, 88, 90, 109, and 111.
- [111] H.-L. Shi, S.-Y. Liu, X.-H. Wang, W.-L. Yang, Z.-Y. Yang, and H. Fan, *Coherence depletion in the Grover quantum search algorithm*, [Physical Review A **95**, 032307 \(2017\)](#), [arXiv:1610.08656](#). Cited on page 41.
- [112] P. W. Shor, *Algorithms for quantum computation: Discrete logarithms and factoring*, in [Proceedings 35th Annual Symposium on Foundations of Computer Science](#) (IEEE Comput. Soc. Press, 1994) pp. 124–134. Cited on page 10.
- [113] U. Skosana and M. Tame, *Demonstration of Shor’s factoring algorithm for $N = 21$ on IBM quantum processors*, [Scientific Reports **11**, 16599 \(2021\)](#). Cited on page 10.
- [114] E. Solano, R. L. de Matos Filho, and N. Zagury, *Deterministic Bell states and measurement of the motional state of two trapped ions*, [Physical Review A **59**, R2539 \(1999\)](#). Cited on page 38.
- [115] A. Sørensen and K. Mølmer, *Quantum Computation with Ions in Thermal Motion*, [Physical Review Letters **82**, 1971 \(1999\)](#), [arXiv:quant-ph/9810039](#). Cited on pages 37, 38, and 74.

Bibliography

- [116] A. Sørensen and K. Mølmer, *Entanglement and quantum computation with ions in thermal motion*, *Physical Review A* **62**, 022311 (2000), [arXiv:quant-ph/0002024](#). Cited on pages 37, 38, 74, and 88.
- [117] 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 (2021). Cited on page 23.
- [118] L. J. Stephenson, D. P. Nadlinger, B. C. Nichol, S. An, P. Drmota, T. G. Ballance, K. Thirumalai, J. F. Goodwin, D. M. Lucas, and C. J. Ballance, *High-Rate, High-Fidelity Entanglement of Qubits Across an Elementary Quantum Network*, *Physical Review Letters* **124**, 110501 (2020). Cited on page 75.
- [119] A. Streltsov, G. Adesso, and M. B. Plenio, Colloquium: *Quantum coherence as a resource*, *Reviews of Modern Physics* **89**, 041003 (2017). Cited on pages 41 and 43.
- [120] A. Streltsov, U. Singh, H. S. Dhar, M. N. Bera, and G. Adesso, *Measuring Quantum Coherence with Entanglement*, *Physical Review Letters* **115**, 020403 (2015). Cited on page 43.
- [121] A. H. Stroud, *Approximate Calculation of Multiple Integrals*, 1st ed. (Prentice-Hall, 1971). Cited on page 84.
- [122] R. T. Sutherland, R. Srinivas, S. C. Burd, D. Leibfried, A. C. Wilson, D. J. Wineland, D. T. C. Allcock, D. H. Slichter, and S. B. Libby, *Versatile laser-free trapped-ion entangling gates*, *New Journal of Physics* **21**, 033033 (2019), [arXiv:1810.08300](#). Cited on page 75.
- [123] S. Szalay, *Multipartite entanglement measures*, *Physical Review A* **92**, 042329 (2015), [arXiv:1503.06071](#). Cited on page 16.
- [124] 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 (2011). Cited on page 25.
- [125] Q. A. Turchette, C. S. Wood, B. E. King, C. J. Myatt, D. Leibfried, W. M. Itano, C. Monroe, and D. J. Wineland, *Deterministic Entanglement of Two Trapped Ions*, *Physical Review Letters* **81**, 3631 (1998). Cited on page 25.
- [126] L. M. K. Vandersypen, M. Steffen, G. Breyta, C. S. Yannoni, M. H. Sherwood, and I. L. Chuang, *Experimental realization of Shor's quantum factoring algorithm using nuclear magnetic resonance*, *Nature* **414**, 883 (2001). Cited on page 10.
- [127] L. Viola and S. Lloyd, *Dynamical suppression of decoherence in two-state quantum systems*, *Physical Review A* **58**, 2733 (1998). Cited on page 74.
- [128] P. Virtanen *et al.*, *SciPy 1.0: Fundamental algorithms for scientific computing in Python*, *Nature Methods* **17**, 261 (2020). Cited on pages 12 and 79.
- [129] K. von Prillwitz, Ł. Rudnicki, and F. Mintert, *Contrast in multipath interference and quantum coherence*, *Physical Review A* **92**, 052114 (2015), [arXiv:1409.1814](#). Cited on pages 44 and 45.
- [130] 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**, 233 (2021). Cited on page 25.
- [131] Y. Wang, M. Um, J. Zhang, S. An, M. Lyu, J.-N. Zhang, L.-M. Duan, D. Yum, and K. Kim, *Single-qubit quantum memory exceeding ten-minute coherence time*, *Nature Photonics* **11**, 646 (2017), [arXiv:1701.04195](#). Cited on page 25.
- [132] A. E. Webb, S. C. Webster, S. Collingbourne, D. Breaud, A. M. Lawrence, S. Weidt, F. Mintert, and W. K. Hensinger, *Resilient Entangling Gates for Trapped Ions*, *Physical Review Letters* **121**, 180501 (2018), [arXiv:1805.07351](#). Cited on pages 75, 79, 88, 89, 90, 109, and 111.
- [133] S. C. Webster, S. Weidt, K. Lake, J. J. McLoughlin, and W. K. Hensinger, *Simple Manipulation of a Microwave Dressed-State Ion Qubit*, *Physical Review Letters* **111**, 140501 (2013). Cited on page 25.

Bibliography

- [134] 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), [arXiv:1603.03384](#). Cited on page 25.
- [135] G. Wendin, *Quantum information processing with superconducting circuits: A review*, *Reports on Progress in Physics* **80**, 106001 (2017), [arXiv:1610.02208](#). Cited on page 11.
- [136] E. B. Wilson, *Probable Inference, the Law of Succession, and Statistical Inference*, *Journal of the American Statistical Association* **22**, 209 (1927). Cited on page 61.
- [137] A. Winter and D. Yang, *Operational Resource Theory of Coherence*, *Physical Review Letters* **116**, 120404 (2016), [arXiv:1506.07975](#). Cited on page 43.
- [138] G. K. Woodgate, *Elementary Atomic Structure*, 2nd ed. (Clarendon Press, Oxford, 1980). Cited on page 30.
- [139] K. Wright *et al.*, *Benchmarking an 11-qubit quantum computer*, *Nature Communications* **10**, 5464 (2019). Cited on page 74.
- [140] G. Zarantonello, H. Hahn, J. Morgner, M. Schulte, A. Bautista-Salvador, R. F. Werner, K. Hammerer, and C. Ospelkaus, *Robust and Resource-Efficient Microwave Near-Field Entangling $^9\text{Be}^+$ Gate*, *Physical Review Letters* **123**, 260503 (2019). Cited on page 75.
- [141] E. J. Zhang *et al.*, *High-fidelity superconducting quantum processors via laser-annealing of transmon qubits*, [arXiv:2012.08475](#) (2020). Cited on page 11.