Device-Independent Key Distribution
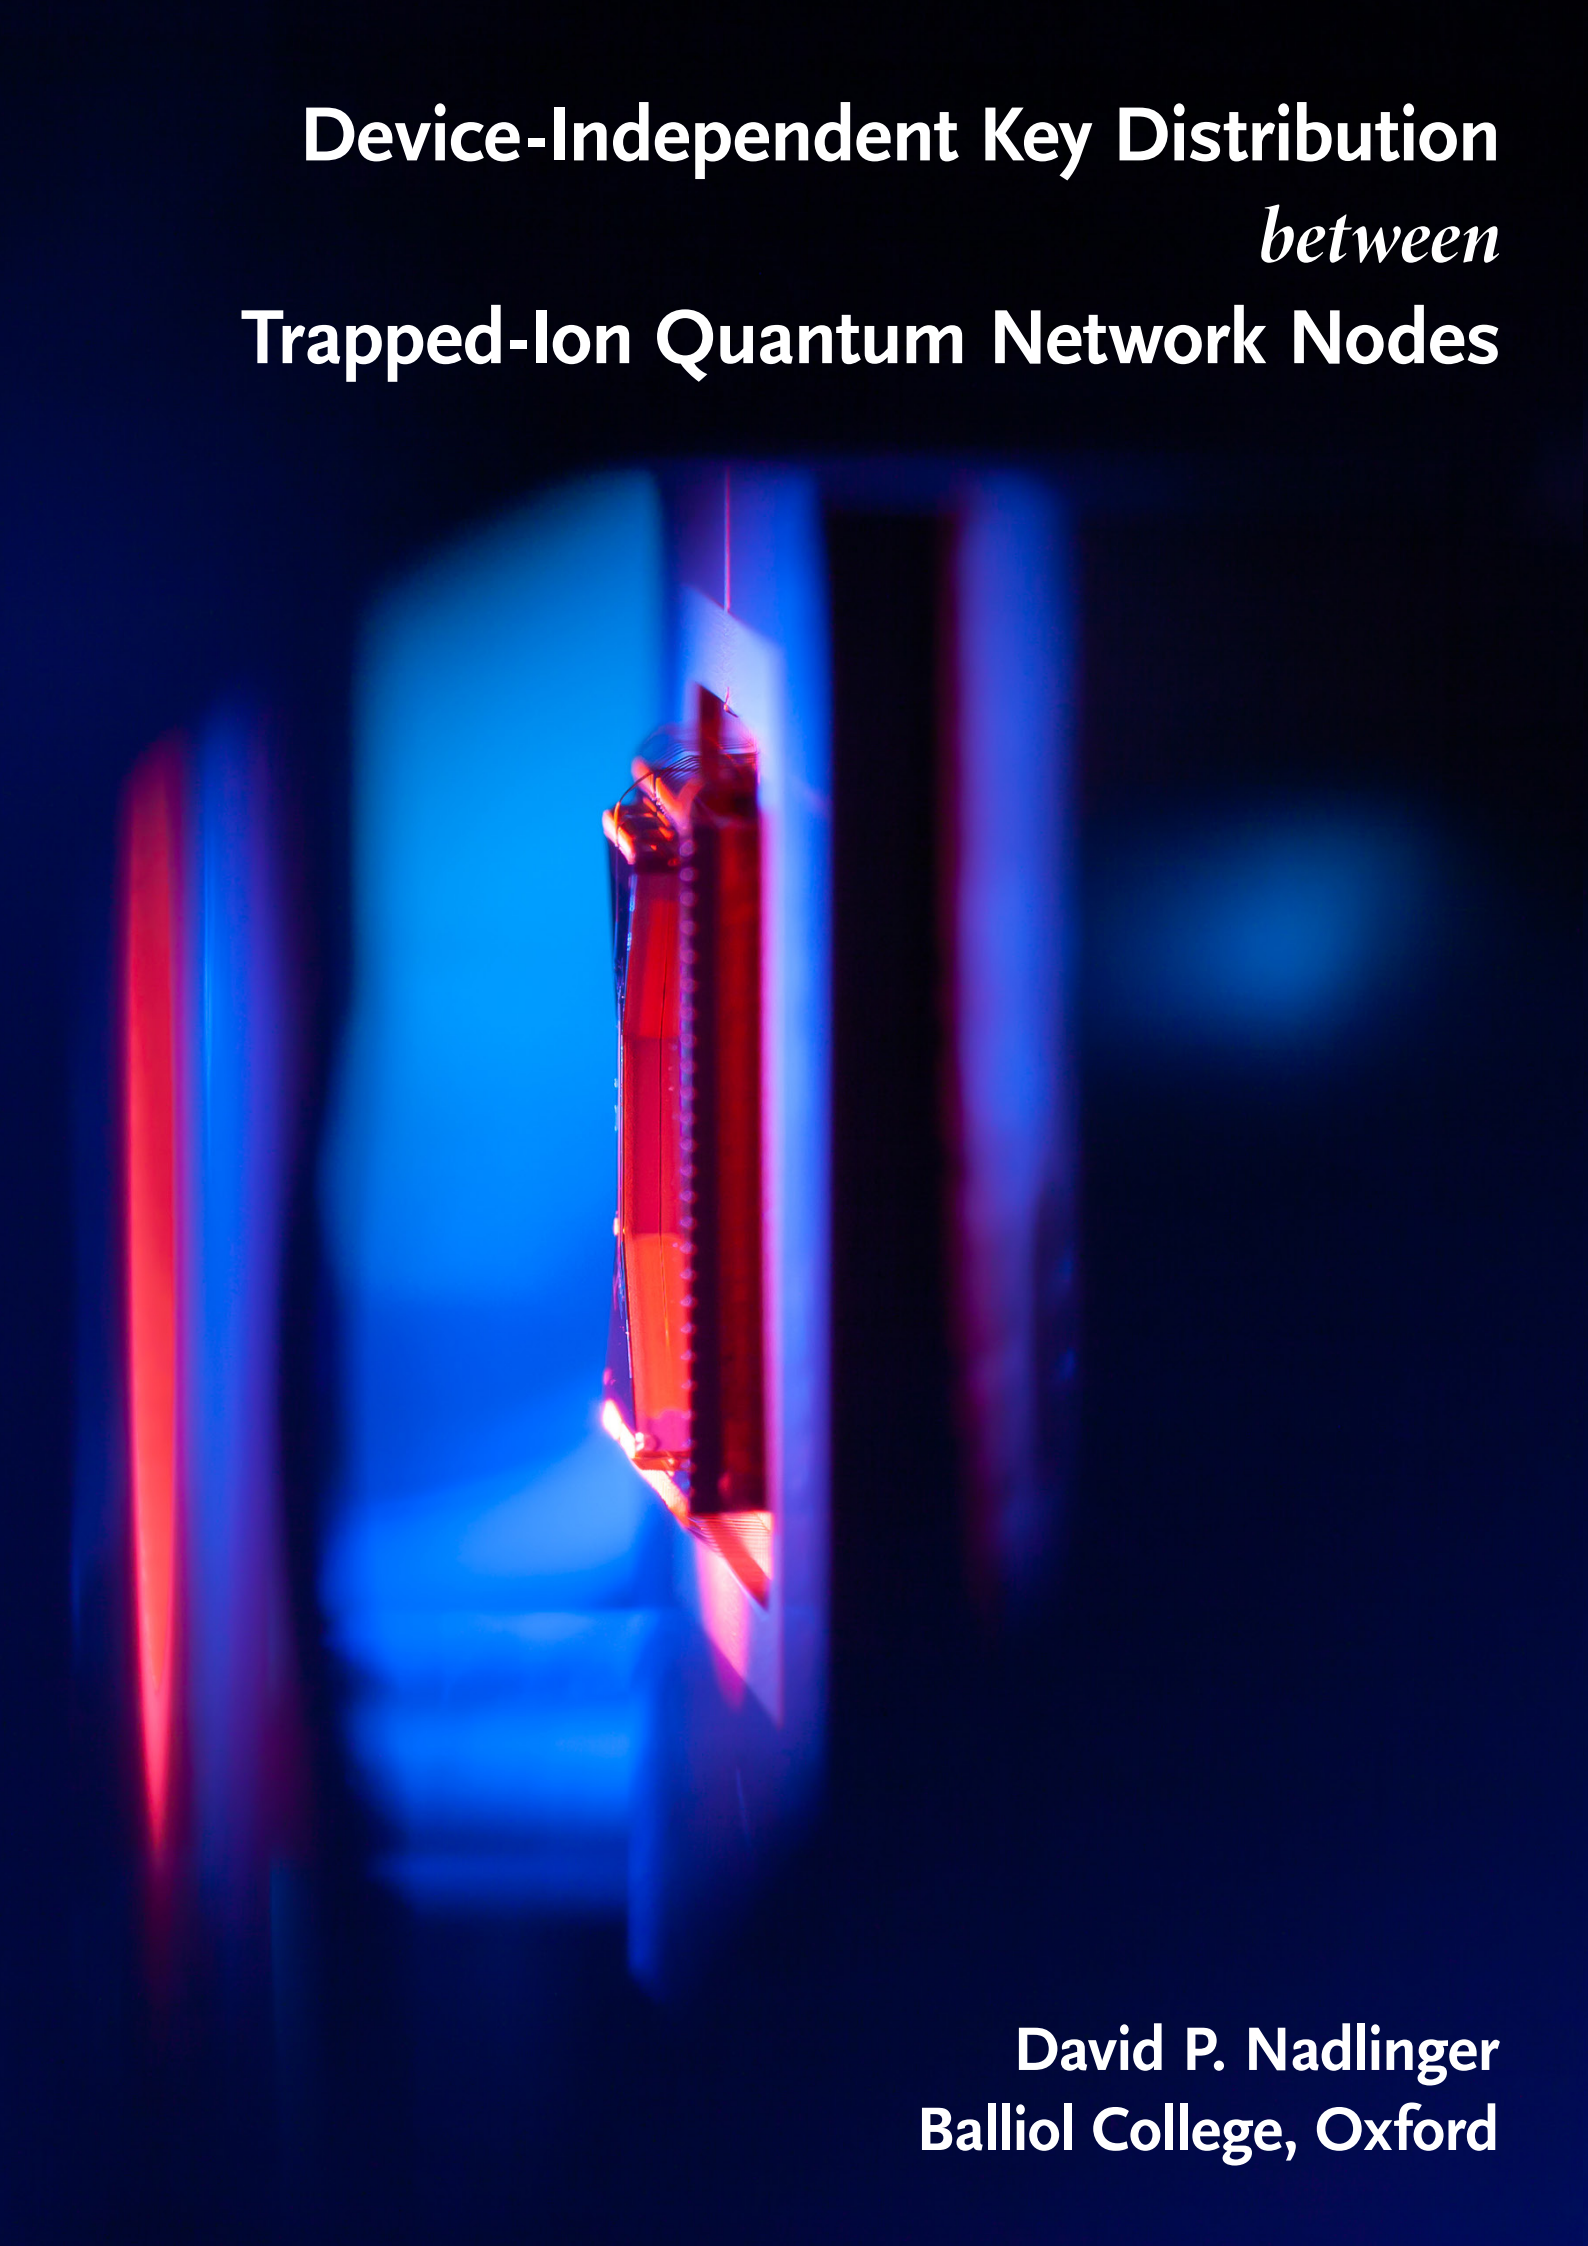
between

Trapped-Ion Quantum Network Nodes

David P. Nadlinger

Balliol College, Oxford

University of Oxford
Department of Physics

Device-Independent Key Distribution between
Trapped-Ion Quantum Network Nodes

David P. Nadlinger

A thesis submitted for the degree of
Doctor of Philosophy

Balliol College
Trinity Term 2022

Device-Independent Key Distribution between Trapped-Ion Quantum Network Nodes
David P. Nadlinger
A thesis submitted for the degree of Doctor of Philosophy
Trinity Term 2022
Balliol College, Oxford

Hybrid quantum systems, combining the advantages of matter-based carriers of quantum information with those of light, have potential applications across many domains of quantum science and technology. In this thesis, we present a high-fidelity, high-rate interface between trapped ions and polarisation-encoded photonic qubits, based on the spontaneous emission of 422 nm photons from ^{88}Sr$^+$, entangled in polarisation with the resulting electronic state of the ion.

We show that photons can be efficiently collected perpendicular to the ambient magnetic field without loss of polarisation purity by exploiting the symmetry properties of single-mode optical fibres, and analyse the impact of a number of common experimental imperfections, including in the heralded entanglement swapping step used to probabilistically generate entanglement between remote ion qubits.

Our experimental platform consists of two ^{88}Sr$^+$–^{43}Ca$^+$ mixed-species quantum network nodes, linked by 2×1.75 m of single-mode optical fibre. We measure an ion–photon entanglement fidelity of 97.7(1) %, generated at an attempt rate of 1 MHz and up to 2.3 % overall collection/detection efficiency. Bell states between remote ^{88}Sr$^+$ ions are generated at a fidelity of 96.0(1) % and rate of $100\,\text{s}^{-1}$. This is the highest fidelity for optically mediated entanglement between distant qubits reported across all matter qubit platforms, and the highest rate among those with fidelities > 70 %.

To compensate stray electric fields that would cause a periodic modulation of the ion position, we introduce a versatile method which relies on the synchronous detection of parametrically excited motion through time-stamped detection of photons scattered during laser cooling. Crucially, only a single laser beam is required to resolve fields in multiple directions; we achieve a stray field sensitivity of $0.1\,\text{Vm}^{-1}/\sqrt{\text{Hz}}$.

Finally, we present the first experimental demonstration of device-independent quantum key distribution, by which two distant parties can share an information-theoretically secure private key even in the presence of an arbitrarily powerful eavesdropper without placing any trust in the quantum behaviour of their devices. This is enabled by a record-high detection-loophole-free CHSH inequality violation of 2.677(6) and low quantum bit error rate of 1.44(2) %, stable across millions of Bell pairs, and an improved security analysis and post-processing pipeline. We implement the complete end-to-end protocol in a realistic setting, allowing Alice and Bob to obtain a 95 884-bit key across 8.5 hours that is secure against the most general quantum attacks.

Our results establish trapped ions as a state-of-the-art platform for photonic entanglement distribution at algorithmically relevant speeds and error rates. The link performance nevertheless remains far from fundamental limits; further improvements are discussed from the perspective of large-scale modular quantum computation as well as from that of long-distance quantum networking applications.

Contents

Acknowledgements

Acknowledgement sections in academic theses seem to have a habit of veering uncomfortably close in genre to that also inhabited by the proverbial Academy Awards acceptance speech: sometimes mildly entertaining, at worst an insufferably lengthy distraction incongruous in tone and pathos with its surroundings, and in either case, less of an opportunity to convey heartfelt gratitude (a duty surely best fulfilled in person!) and more of an exercise in avoiding to snub any potential readers who might expect to find their name among those mentioned.

Be that as it may, I would like to acknowledge the collaborative origins of the work presented here. Experimental physics of this kind, right at the boundary of technical feasibility, would be all but impossible without collaboration (the organisation of which is not in the least made easier by the fact that the vast amount of necessary engineering-adjacent work is by itself seldom valued in the wider academic context). At the end of my undergraduate degree at ETH Zürich, I made the decision to come to Oxford – not an easy one, as I had gotten an excellent start in ion trapping in Prof. Jonathan Home's ETH group. In Oxford, I had the opportunity to see through the local ion trap quantum networking effort from its humble beginnings, around the time I obtained the first proof-of-principle ion–photon entanglement results as a visiting MSc student, to where we now have the first (and second, and third, …) glossy results from our own two-node quantum network in hand, with more surely to come.

By name, let me extend my thanks to some of my colleagues who are or were part of the "core team" who together tackled the daunting challenge of building from scratch and maintaining (and rebuilding) not one, but two dual-species ion trap experiments: Chris Ballance, Laurent Stephenson, Tim Ballance, Joe Goodwin, Bethan Nichol, Peter Drmota, Dougal Main, Gabriel Araneda Machuca, and Raghavendra Srinivas (all of whom are either Drs., or will surely be before long). Separately, I must thank Prof. David Lucas not just for his role in the particular projects discussed here, but also for his steadfast administrative support; as his students, we were always free to get on with our work in the lab without needing to worry about being caught off guard by formal constraints or financial matters.

I would also like to thank all other members past and present of the Oxford Ion Trap Quantum Computing group for the opportunity to collaborate on a number of exciting projects and publications during all my years here, and those who offered to read earlier

1   *Introduction*

Quantum information science and technology are the result of one of the most extraordinary intellectual journeys over the last century. Its origins are in turn-of-the-century Europe, where quantum mechanics had just been invented; back then firmly part of fundamental physics, to settle questions such as that of the nature of light and the rules that govern nature on atomic scales, and bringing with it a number of new concepts, such as superposition and entanglement. Onto this scientific revolution soon followed a technological one: the deepened understanding of chemistry and material science, most prominently that of semiconductors, led to a number of breakthrough technological innovations in the 1950s and 1960s, such as lasers and transistors, which would later form the bedrock of the modern information age.

Now, in the twenty-first century, we find ourselves right in the middle of a second quantum revolution [DM03]. No longer are we restricted to only observing the classical shadows of the underlying quantum phenomena in the behaviour of bulk materials; we can now actively manipulate single, isolated quantum systems. Together with a greatly matured understanding of the structure of quantum theory and its ramifications for information processing, this has given rise to a number of endeavours in various stages along the path from conceptual proposal to real-world technology. The efficient simulation of other quantum systems is one of the oldest proposals for the use of tailor-made, engineered quantum systems, going back to a landmark speech by R. Feynman [Fey82], with many potential applications across the physical sciences wherever quantum effects would make the modelling on classical computers intractable [GAN14]. Quantum computing as a form of digital information processing fundamentally distinct from classical Turing machines goes back to D. Deutsch [Deu85]; few have doubted its potential utility ever since P. Shor's discovery of an efficient quantum algorithm for the factorisation of large numbers [Sho94]. Much attention across academia and industry is currently focussed on both the development of useful quantum algorithms as well as the hardware to execute them on, in particular towards the development of fault-tolerant quantum computing, where error correction codes [Sho95; Ste96] mitigate decoherence and other hardware imperfections to enable large-scale computations [Pre18]. Another prominent technological use of quantum effects is to augment the precision of sensors [DRC17]. For geographically distributed quantum networks, a long list of applications has been proposed too, ranging from secure
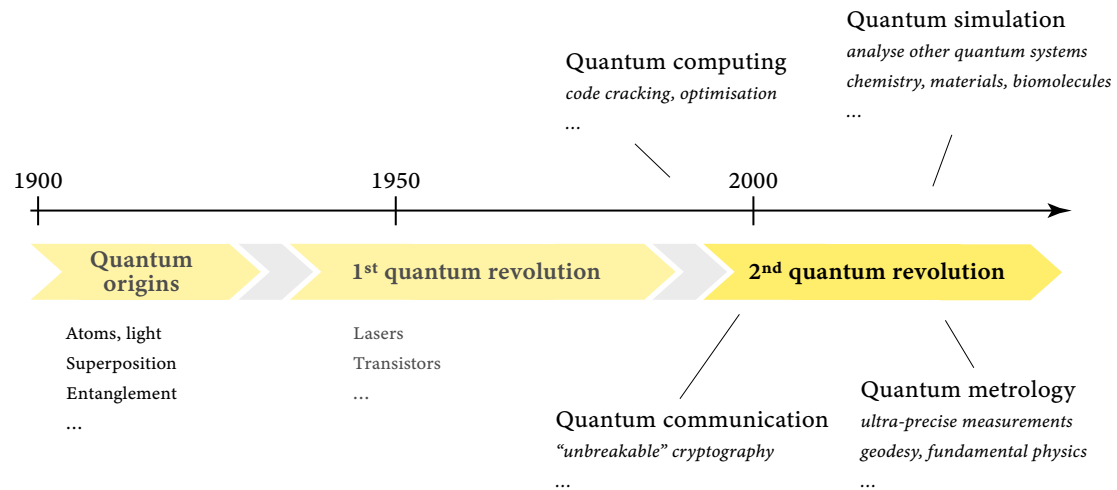
Figure 1.1: Cartoon timeline of (some of the) developments in quantum physics in the 20th century, illustrating its progression from a study of the fundamental laws of nature, to an indirect substrate for technical developments such as semiconductors, until finally isolated quantum systems themselves gain a role in new kinds of information processing.

communication to enhanced metrology to fundamental physics research – somewhat aspirationally referred to as "the quantum internet" [Kim08; WEH18].

1.1   *Trapped ions and photons*

Just as the same classical information can be stored bit-by-bit in the magnetisation pattern of a platter in a computer hard drive, held in working memory as charges on small capacitors, or transmitted across long distances as laser pulses in optical fibres, the abstract concept of a bit of quantum information, a qubit, can similarly be realised in a number of different physical forms. Large-scale quantum information processing has a diverse set of requirements, and no single physical implementation of quantum information equally excels across all of them.

Trapped atomic ions are a promising candidate platform for qubits in large-scale quantum computing; their state can be accurately initialised and read out [Har+14; Chr+20], they are impeccable memories [Sep+19; Wan+21], and can be readily made to interact with microwave [Har+16; Sri+21] or laser fields [Bal+16; Gae+16; Cla+21] to implement logic operations with record-high fidelity. Nevertheless, their dependence on an ultra-high vacuum environment typically restricts them to a single, immobile vacuum chamber (though they can be "shuttled" around with a trap apparatus without disturbing the stored quantum state on account of their electric charge). Conversely, photons excel in communication applications; they naturally travel at the speed of light and can be guided in regular optical fibres while maintaining a reasonable amount of quantum coherence over tens of kilometres. However, they are – at the single-photon level – difficult to couple together

deterministically to implement logic operations, difficult to generate "on demand", difficult to store in place, and the measurement of their state is hampered by finite detector quantum efficiencies.

Hybrid systems that combine the complementary strengths of trapped ions (or other matter qubit systems) and photons are thus attractive in a number of scenarios. For this, an interface to convert information between different qubit types is necessary that is swift and generates highly coherent states, so as neither to stall a computation nor corrupt its results. Photon-mediated entanglement of remote matter qubits has been studied theoretically since around the year 2000 [Cir+97; Dua+01], and remote entanglement has been demonstrated in a number of different qubit types since; neglecting pure quantum memories (where a photonic excitation is stored for later retrieval, but cannot directly manipulated), these include trapped ions [Moe+07], neutral atoms [Rit+12], nitrogen-vacancy centres in diamond [Ber+13], solid-state quantum dots [Del+16], and micro-mechanical oscillators [Rie+18].

In the context of trapped-ion quantum information processing, one particularly interesting use case for such hybrid systems is as one of several tools to address a major challenge in scaling up ion registers to large numbers of qubits: "spectral crowding", the increasing difficulty in controlling the spectrum of vibrational modes in a single linear chain of ions as its length increases [MK13]. One strategy is to keep multiple strings of ions in a single trap with segmented electrodes [Win+98; KMW02]; splitting, moving and recombining chains as necessary to facilitate the desired interactions through local Coulomb interactions. Optical links could augment this by providing long-range connectivity, or by reducing the technical complexity by allowing a large-scale processor to be constructed from multiple separate modules connected by a flexible fibre-optical fabric, each perhaps only containing a relatively small number of ions [Mon+14; NFB14].

1.2   An elementary two-node quantum network

In this work, we take for granted the mature toolbox of experimental techniques for trapped-ion quantum information processing and concentrate on the development of a free-space photonic interconnect. The interface is based on the spontaneous emission of, in this particular instance, $422\,\mathrm{nm}$ photons from $^{88}\mathrm{Sr}^{+}$, entangled in polarisation with the resulting electronic state of the ion. We show that photons can be collected in the (technologically favourable) direction perpendicular to the magnetic field using high-numerical-aperture lenses while retaining unit fidelity of the resulting entangled state; for this, we exploit the symmetry properties of single-mode optical fibres to filter out contributions otherwise leading to deleterious polarisation mixing effects. As such, the practically attainable performance is necessarily limited by technical imperfections. We present a framework for the unified treatment of such sources of errors in the ion–photon

entanglement generation process in each ion-trap network node, as well as the photonic Bell-basis measurement on pairs of such photons that implements the heralded generation of remote ion–ion entanglement between nodes through entanglement swapping. We quantitatively catalogue a number of likely sources of error, finding that generation of Bell states between remote ions with fidelities in excess of 99 % should be feasible assuming realistic amounts of imperfection.

We practically implement this remote entanglement generation scheme in an elementary network of two ^{43}Ca$^+$–^{88}Sr$^+$ quantum network nodes linked by 2×1.75 m of optical fibre; the ability to manipulate multiple strings of ions in the microfabricated surface electrode traps employed and the availability of ^{43}Ca$^+$ as application qubits with long coherence times will allow the demonstration of a number of few-qubit quantum networking protocols. Here, we focus on experiments with a single ^{88}Sr$^+$ interface ion per node; characterising the quantum correlations between ions and photon polarisation to qualitatively validate the theoretical treatment, and observing ion–photon entanglement with a fidelity of 97.7(1) %, generated at a rate of ≈ 23 ms^{-1}. We similarly analyse the remote ion–ion entanglement using state tomography, observing a Bell-state fidelity of $\mathcal{F} = 96.0(1)$ % at a rate of 100 s^{-1} (up to 182 s^{-1}). This vastly improves on the performance of previous remote matter qubit entanglement experiments.

One requirement for achieving this was the minimisation of stray electric fields in the trap that would cause a periodic modulation of the ion position through interaction with the radio-frequency trap drive, and thus the wavepacket of the entangled photons emitted. We present a precise method for this which relies on parametric modulation of the trapping potential, while recording and de-modulating the arrival times of photons scattered during laser cooling to recover the thusly excited motion. Only a single laser beam is required to resolve fields in multiple directions, and we achieve a stray field sensitivity of 0.1 Vm$^{-1}/\sqrt{\text{Hz}}$ and uncertainty floor of 0.015 Vm^{-1}, which favourably compares with other technically more complex or less robust methods.

Finally, we report on the first experimental demonstration of device-independent quantum key distribution (DIQKD): a method for two distant parties to share secret encryption keys even in the presence of an arbitrarily powerful eavesdropper. This information-theoretic level of security is fundamentally impossible using classical communication; furthermore – and perhaps surprisingly – by relying on the properties of quantum entanglement, it is possible to guarantee security even if the parties do not trust their quantum devices to behave as intended. Though proposed already in 1991 by A. Ekert [Eke91], a demonstration has thus far remained out of reach due to its stringent requirements on the quantum hardware. Making use of our remote entanglement link, we obtain a record-high detection-loophole-free Clauser–Horne–Shimony–Holt (CHSH) inequality violation of $S = 2.677(6)$ and low quantum bit error rate of $Q = 1.44(2)$ %, stable across millions of Bell pairs acquired at an overall average rate of 63 s^{-1} (reduced

from the raw rate of $129\,\mathrm{s}^{-1}$ due to frequent ion reloading events due to vacuum issues). Combining this with improvements to the theoretical security proof and post-processing pipeline, we obtain a 95 884-bit key with fully device-independent security (soundness error $\epsilon_{\mathrm{snd}} = 10^{-10}$). Reaching this "ultimate limit to privacy" [ER14] not only represents a milestone in quantum communication, but also unambiguously demonstrates that the link is stable enough to be used as a component in more complex experiments.

1.3 Thesis layout

The remainder of this thesis is structured as follows:

Chapter 2 reviews the basic notions of the theoretical description of atomic ions in radio-frequency traps, of photons, and of their interaction; particularly regarding the generation of entanglement through spontaneous emission.

Chapter 3 gives a brief overview of the experimental apparatus making up the elementary quantum network: two $^{88}\mathrm{Sr}^+/^{43}\mathrm{Ca}^+$ ion-trap nodes and the photonic link connecting them.

Chapter 4 considers in detail geometric aspects of the mechanism for entanglement generation between a $^{88}\mathrm{Sr}^+$ ion and the polarisation of a spontaneous decay photon, and the subsequent establishment of entanglement between two remote $^{88}\mathrm{Sr}^+$ ions through heralded entanglement swapping. We argue that coupling of the emission into a single-mode fibre avoids errors in the generated state for reasons of symmetry, and catalogue the effect of a number of relevant experimental imperfections.

Chapter 5 introduces a versatile method for stray field compensation in Paul traps through synchronous detection of parametrically excited motion. We provide an analytical description, showing that the ion response is linear in a possible stray field, which we verify in the experiment, and discuss an algorithm for multi-dimensional stray field compensation. A journal manuscript based on the contents of this chapter is in preparation [Nad+21].

Chapter 6 describes the characterisation of the entanglement generation primitives – between a single ion and a single photon, and after entanglement swapping, between two remote ions – in the experiment, and compares the achieved performance to a comprehensive literature survey. The results of this analysis (though for an earlier set of experimental data) were partially published in ref. [Ste+20].

Chapter 7 describes the use of this quantum network for the first demonstration of device-independent quantum key distribution, covering both the theoretical background and the details of the experimental implementation. This result builds on joint work with colleagues in quantum information theory and classical coding theory, and has been published in ref. [Nad+22] (though the description here is mine).

Chapter 8 concludes this thesis with an outlook towards further improvements to the remote entanglement generation performance, an increased distance between the ion trap nodes, and the next steps towards more complex quantum networking applications.

Not described here is a recent demonstration using the same apparatus for the entanglement-enhanced comparison of atomic clocks [Nic+22] (see the doctoral thesis of B. Nichol [Nic22]), and neither are separate characterisations of $^{88}\text{Sr}^+$–$^{43}\text{Ca}^+$ quantum logic gates [Hug+20] and a $^{43}\text{Ca}^+$ quantum memory [Sep+19] in different experimental apparatus (see theses of K. Thirumalai [Thi19] and A. Hughes [Hug21]).

2 Background

2.1 Quantum formalism

The quantum-theoretical formalism used throughout this thesis is almost without exception[1] the canonical one of states and density operators described in any number of contemporary textbooks (e.g. ref. [NC11]). We thus only give a brief summary here to establish the notational choices and to give a few useful results.

A pure quantum state is an element $|\psi\rangle$ of a complex Hilbert space $\mathcal{H}$ with norm $\|\psi\|^2 = \langle\psi|\psi\rangle = 1$. To be able to express statistical mixtures (ensembles) of multiple states, we use density operators; elements of the space $\mathcal{S}_=(\mathcal{H})$ of bounded, self-adjoint operators $\rho \in \mathrm{End}(\mathcal{H})$ that are positive ($\rho \geq 0$) and have unit trace ($\mathrm{tr}\,\rho = 1$). Pure states are equivalent to projectors $\rho = |\psi\rangle\langle\psi|$. To conveniently capture the notion of a state of interest only arising with a certain probability, we will also allow sub-normalised states $\rho \in \mathcal{S}_\leq(\mathcal{H})$, for which the trace condition is modified to $\mathrm{tr}\,\rho \leq 1$. Some authors prefer to adorn the symbols for all operators with a hat (e.g. $\hat{\rho}$); here, we do not follow this convention and instead use $\hat{\cdot}$ to refer to vectors of unit length.

A general quantum measurement is given by some alphabet I_M labelling the outcomes (events) $m \in I_M$, and an associated set of self-adjoint operators M_m, which is complete in the sense $\sum_{m \in I_M} M_m^\dagger M_m = \mathbb{1}$. For a given state ρ, the probability to observe each outcome is $P_m(\rho) = \mathrm{tr}\left(M_m^\dagger M_m \rho\right)$. We will write $\rho_{\wedge m} \coloneqq M_m \rho M_m^\dagger$ for the sub-normalised state conditioned on the outcome m ($\mathrm{tr}\,\rho_{\wedge m} = P_m(\rho)$), and $\rho_{|m} \coloneqq \rho_{\wedge m}/P_m(\rho)$ for the normalised state after having observed the outcome m. If only the probabilities $P_m(\rho)$ are of interest, but not the post-measurement state, it is enough to consider the positive operators $E_m \coloneqq M_m^\dagger M_m$, as $P_m(\rho) = \mathrm{tr}\left(E_m \rho\right)$. The entire set $\{E_m\}_{m \in I_m}$ is known as a positive operator-valued measure, or POVM, the individual operators E_m as POVM elements.

Two systems A and B are composed by taking the tensor product of their Hilbert spaces $\mathcal{H}_A$ and $\mathcal{H}_B$, $\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B$. Where it aids clarity, we will write ρ_{AB} for the full, joint

---

[1] The one exception are the string diagrams used to illustrate entanglement-based protocols (e.g. fig. 2.4), which can be formalised using category theory (see e.g. [Coe10]). While this perspective is almost uniquely well-suited for the intuitive understanding of protocols such as entanglement swapping, we do not develop it further here due to space constraints. The interested reader is referred to a recent textbook by Kissinger and Coecke [CK18].
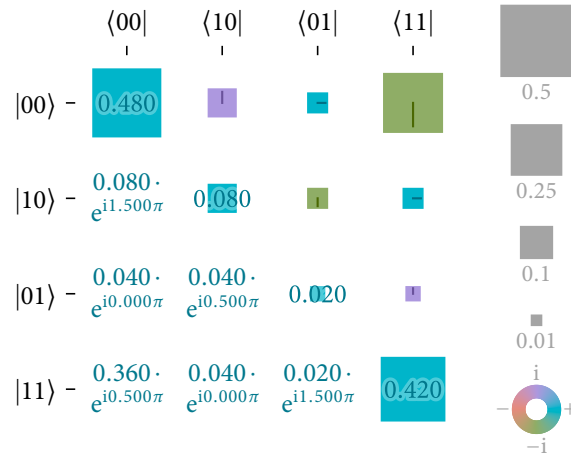
Figure 2.1: Graphical representation of two-qubit density matrices employed throughout this thesis, illustrated by example of the arbitrarily chosen state $\rho = \frac{4}{5}|\psi\rangle\langle\psi| + \frac{1}{5}|\phi\rangle\langle\phi|$, with $|\psi\rangle = \frac{1}{\sqrt{2}}(|00\rangle + i|11\rangle)$ and $|\phi\rangle = (|0\rangle - i|1\rangle)/\sqrt{2} \otimes (2|0\rangle + |1\rangle)/\sqrt{5}$. The area of each square in the upper-right half is proportional to the magnitude of the corresponding matrix element; its colour represents the complex phase according to the key shown (which is also indicated by small "clock hands" for monochrome viewing). The numbers in the lower-left half indicate the same information in numerical form. These two forms are equivalent, as physical density matrices are Hermitian (which applies also to the maximum-likelihood reconstructions in later chapters of this thesis).

state of the systems, and e.g. $\rho_A = \operatorname{tr}_B \rho_{AB}$ for the reduced state of subsystem A on its own. If $\{|b\rangle\}_{b \in I_B}$ is an orthonormal basis for B, the partial trace tr_B can be evaluated as

$$\rho_A = \operatorname{tr}_B \rho_{AB} = \sum_{b \in I_B} \langle b|\rho_{AB}|b\rangle. \tag{2.1}$$

The von-Neumann entropy of a quantum state is

$$H(\rho) := -\operatorname{tr}(\rho \log_2 \rho), \tag{2.2}$$

which, for a two-dimensional classical state $\rho = p|0\rangle\langle 0| + (1-p)|1\rangle\langle 1|$ coincides with the usual binary entropy function from classical information theory,

$$h(p) := -p \log_2 p - (1-p) \log_2(1-p). \tag{2.3}$$

For a state ρ_{AB} on subsystems A and B, an alternative notation is $H(A, B)_\rho := H(\rho_{AB})$, $H(A)_\rho := H(\rho_A)$, etc. The conditional entropy of A given B is

$$H(A|B)_\rho := H(A, B)_\rho - H(B)_\rho, \tag{2.4}$$

and can be negative for entangled states.

2.1.1   One and two qubits

In particular, for a qubit, $\mathcal{H} = \mathbb{C}^2$, and

$$|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle \in \mathbb{C}^2, \quad |\alpha|^2 + |\beta|^2 = 1, \tag{2.5}$$

where $\{|0\rangle, |1\rangle\}$ is any particular choice of orthonormal basis, referred to as the computational basis. The Pauli matrices

$$\sigma_1 := X := \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \sigma_2 := Y := \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \quad \sigma_3 := Z := \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \tag{2.6}$$

form a basis for all traceless Hermitian 2×2 matrices as a real vector space, and thus, together with the identity matrix, all Hermitian 2×2 matrices (so, in particular, the single-qubit density matrices). Scaled by $\frac{1}{\sqrt{2}}$, this basis is orthonormal under the Hilbert–Schmidt product $\langle A, B \rangle = \mathrm{tr}(A^\dagger B)$. For normalised states ρ, coefficients $r_i = \mathrm{tr}(\sigma_i \rho)$ in this basis define a three-dimensional real vector $\boldsymbol{r}$, known as the Bloch vector, which uniquely identifies single-qubit density matrices. $|\boldsymbol{r}| \leq 1$ follows from $\mathrm{tr}\,\rho = 1$. The pure states, up to a global phase in $\mathrm{U}(1) \in \mathbb{C}$, saturate this bound and lie on the unit sphere, referred to as the Bloch sphere. There is a surjective homomorphism between single-qubit operations $U \in \mathrm{U}(2)$ and rotations of the Bloch sphere in $\mathrm{SO}(3)$, such that the Bloch sphere lends itself to the intuitive understanding of one-qubit quantum operations[2].

For two qubits, $\mathcal{H} = \mathbb{C}^2 \otimes \mathbb{C}^2$, a prominent alternative to the computational basis $|00\rangle = |0\rangle \otimes |0\rangle$, $|01\rangle = |0\rangle \otimes |1\rangle$, ... is the Bell basis given by the four Bell states

$$\left|\Phi^\pm\right\rangle = \frac{1}{\sqrt{2}}\left(|00\rangle \pm |11\rangle\right), \quad \left|\Psi^\pm\right\rangle = \frac{1}{\sqrt{2}}\left(|01\rangle \pm |10\rangle\right). \tag{2.7}$$

The Bell states are maximally entangled (cf. §7.1.2); as such, they can be transformed into each other using only local unitary operations, i.e. operations that act on only of the qubits.

2.1.2   Distance metrics

Given two density operators $\rho, \sigma \in \mathcal{S}(\mathcal{H})$, there are several widely used ways of defining a notion of distance between them. The trace distance $\delta(\rho, \sigma) : \mathcal{S}(\mathcal{H})^2 \to [0, 1]$ is a metric on the space of density operators, and between two density operators ρ and σ can be defined in terms of the trace norm $\|S\|_1 = \mathrm{tr}\,\sqrt{S^\dagger S}$ as

$$\delta(\rho, \sigma) := \frac{1}{2}\|\rho - \sigma\|_1. \tag{2.8}$$

---

[2] The homomorphism is in fact a double cover due to the possibility of a global phase, but otherwise unique. I have previously given a concise, self-contained summary of the standard results concerning the representation of unitary operations as Bloch sphere rotations in ref. [Nad16, appendix A].

It generalises the classical notion of statistical distance, and has an alternative interpretation as

$$\delta(\rho, \sigma) = \max_{\substack{\Lambda \in \text{End}(\mathcal{H}) \\ 0 \leq \Lambda \leq 1}} \text{tr}\left(\Lambda(\rho - \sigma)\right), \tag{2.9}$$

i.e. $\frac{1}{2}(1 + \delta(\rho, \sigma))$ is the limit to the probability of correctly identifying a system prepared in either ρ or σ uniformly at random with a single measurement (saturated by the optimal POVM element Λ).

The fidelity $F(\rho, \sigma)$ is a generalisation of the state overlap for pure states,

$$F(|\psi\rangle\langle\psi|, |\phi\rangle\langle\phi|) = |\langle\psi|\phi\rangle|^2, \tag{2.10}$$

and can be expressed as[3]

$$F(\rho, \sigma) = \left(\|\sqrt{\rho}\sqrt{\sigma}\|_1\right)^2. \tag{2.11}$$

2.1.3   Quantifying entanglement

Maximally entangled states, such as the four Bell states, are an important resource in many quantum information processing scenarios (such as the entanglement-based quantum key distribution protocol discussed in chapter 7). The precise form of the generated state is typically not important, as the local unitary operations necessary to transform one maximally entangled state into another are typically much easier to implement with low error rates than the entangling operations (e.g. for trapped ions, single-qubit gates with an error rate of $\sim 1 \times 10^{-6}$ have been demonstrated [Har+16]). When evaluating remote entanglement generation schemes – especially in the context of theoretical modelling[4] – it is thus convenient to find a metric that quantifies the "amount" of entanglement in the generated state irrespective of the particular form of the state.

---

[3]Another convention exists where the fidelity is taken to be the square root of F as defined here. There does not seem to be a clear preference for either (cf. refs. [GLN05; NC11]), though most of the experimental literature uses the "squared" definition that we also adopt.

[4]An entanglement measure that is independent of the specific form of any maximally entangled state can be especially helpful when systematically studying the effect of experimental parameters or imperfections, as it allows one to cleanly separate benign effects that can be easily counteracted using local unitaries from those that actually degrade the entangled state, without having to at every step explicitly take account of this distinction. As a concrete illustration, consider the case of birefringence in the optical fibres used for remote entanglement generation discussed in §4.5.2/fig. 4.24. There, the effect ranges from being completely correctable to unavoidably detrimental based on a parameter of the system (the normalised qubit frequency splitting), which is elegantly captured by the target-state-independent entanglement measure chosen. Note, however, that to make use of the resulting entanglement in a wider circuit or protocol, the actual state generated must usually be known. Consequently, the generated state must in practice be stable enough to allow for calibration of the necessary local corrections ahead of time, which can then be applied during protocol execution.

The most widely quoted quantity in experimental literature discussing the generation of particular states is the fidelity F. We can define the fidelity to the closest maximally entangled state as

$$\mathcal{F}(\rho_{AB}) \coloneqq \max_{\substack{U=U_A\otimes U_B \\ U_A, U_B \in \mathrm{U}(2)}} F(\rho_{AB}, U\,|\Phi^+\rangle\langle\Phi^+|\,U^\dagger). \tag{2.12}$$

$\mathcal{F}$ is also known as the fully entangled fraction (or, slightly vaguely, as Bell-state fidelity) and can be explicitly expressed in the form of a singular value problem [Bad+00]. We also define the Bell-state error as its complement,

$$\mathcal{E}(\rho_{AB}) \coloneqq 1 - \mathcal{F}(\rho_{AB}). \tag{2.13}$$

In preparation for chapter 4, where we will study protocols generating the antisymmetric Bell states $|\Psi^\pm\rangle$, consider a two-qubit density matrix ρ close to these states,

$$\rho = \frac{1}{2}\begin{pmatrix} t+s & 0 & 0 & 0 \\ 0 & 1-t+r & \overline{a} & 0 \\ 0 & a & 1-t-r & 0 \\ 0 & 0 & 0 & t-s \end{pmatrix}, \tag{2.14}$$

where $0 \le t \le \frac{1}{2}$ and $r, s \in \mathbb{R}$, $a \in \mathbb{C}$ are arbitrary (though constrained by ρ being a density matrix). The maximally entangled state that maximises the overlap with such a state ρ (such that $\mathcal{F}(\rho) = \langle\psi|\rho|\psi\rangle$) is $|\psi\rangle = \frac{1}{\sqrt{2}}\left(|01\rangle + e^{i\arg a}\,|10\rangle\right)$, and the fully entangled fraction is

$$\mathcal{F}(\rho) = \frac{1}{2}\left(1 - t + |a|\right). \tag{2.15}$$

Another state-agnostic quantity is the entanglement of formation E_F [Ben+96], defined as

$$E_F(\rho) = \min\left\{ \sum_k p_k E(|\psi_k\rangle) \,\middle|\, \rho = \sum_k p_k |\psi_k\rangle\langle\psi_k| \right\}, \tag{2.16}$$

where $E(|\psi\rangle) = H(\mathrm{tr}_A(|\psi\rangle\langle\psi|)) = H(\mathrm{tr}_B(|\psi\rangle\langle\psi|))$ is the entanglement entropy. Loosely speaking, E_F thus quantifies how much entanglement is on average necessary to generate ρ. For two-qubit states, a closed-form expression for E_F is given in ref. [Woo98]. A further entanglement measure of interest is the logarithmic negativity [VW02] defined as

$$\mathcal{N}(\rho_{AB}) = \log_2\|\rho_{AB}^\Gamma\|_1, \tag{2.17}$$

where $\cdot^\Gamma$ is the partial transpose. $\mathcal{N}$ is non-increasing under local operations and classical communication (LOCC), and upper-bounds the distillable entanglement [Hor+09] (which is not efficiently computable).
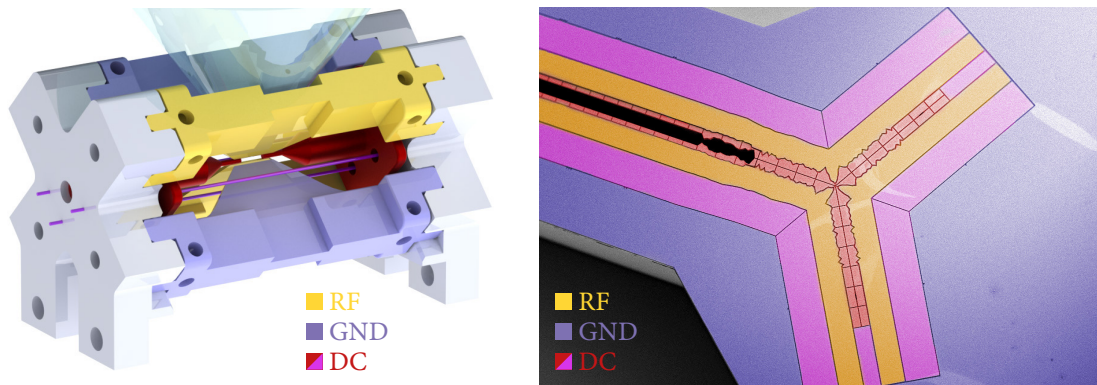
Figure 2.2: Examples of Paul traps for the confinement of electrically charged particles. Radio-frequency (RF) voltages are applied to some electrodes, whereas others are held at ground (GND) or static potentials (DC), creating in the centre an effective potential that is confining in every directions. *Left:* A classic, macroscopic "blade" trap used in another experiment in our laboratory. The ion-electrode distance is 500 μm. *Right:* A surface electrode trap (Sandia *HOA2*), where the trap electrode structure has been "flattened" into a plane, creating a pseudopotential minimum 70 μm above the trap surface. Complex electrode designs can be produced through microfabrication techniques; the segmented DC electrodes allow for many independent, movable potential wells to be created along the trap axis, and multiple linear regions to be connected by junctions. A slot is also visible, which allows optical access from the back side (original scanning electron microscope picture courtesy of Sandia National Laboratories; manually colorised).

2.2   Trapped ions

2.2.1   Charged particles in radio-frequency traps

For our purposes throughout this thesis, trapped ions might well be idealised as abstract n-level systems; coupled to the electromagnetic field, but otherwise suspended at a fixed point in space[5]. The details of the trap mechanism are, in principle, of little interest, as long as it allows us to "pin" an ion at a certain position in space. As the electric field is divergence-free in a source-free region (i.e., the potential satisfies Laplace's equation), it is not possible to achieve this through electrostatic means. However, exploiting the non-zero mass of the particle, it is possible to obtain an effective confining potential in all spatial directions as the time average of a divergence-free quadrupolar "saddle" potential that oscillates quickly between confining and anti-confining along two orthogonal directions.

This is the operating principle of the Paul trap, sometimes also known as radio-frequency trap (referring to the range of oscillation frequencies typically chosen for the confinement of atoms). Figure 2.2 shows two examples: An oscillating quadrupole potential, generated by elongated electrodes to which RF voltages are applied relative to the ground potential, confines the particles to one axis (radial confinement); additional static electrodes generate a potential well also along that axis (axial confinement). Such traps are widely used across different fields, including for the confinement of single

---

[5]The only situation in which this simple picture is not appropriate is going to be the treatment of position- or velocity-dependent imperfections. Chapter 5 describes a technique aimed at reducing trap-induced motion to restore validity of this simple picture as far as possible.

atoms [Lei+03], electrons [Mat+21], molecules [Loh+13], nanoparticles [Ald+16] and microdiamonds [Del+20].

To describe the dynamics of a point-like particle of mass m and charge Ze in a Paul trap with drive frequency Ω_{rf}, it is convenient to expand both the static potential Φ_{dc} and the oscillating potential $\Phi_{\mathrm{rf}} \cos(\Omega_{\mathrm{rf}}t)$ around the equilibrium position, which we assume, without loss of generality, to be in the coordinate origin. In the absence of stray fields, the respective gradients $\nabla\Phi_{\mathrm{dc}}$ and $\nabla\Phi_{\mathrm{rf}}$ vanish at this point, and the equations of motion for particle position $\boldsymbol{x} = (x_i)_{i=1,2,3}$ are

$$\ddot{\boldsymbol{x}} = -\frac{Ze}{2m}\left(\boldsymbol{H}_{\Phi_{\mathrm{dc}}} + \boldsymbol{H}_{\Phi_{\mathrm{rf}}} \cos(\Omega_{\mathrm{rf}}t)\right)\boldsymbol{x} + \mathrm{O}(|\boldsymbol{x}|^2), \qquad (2.18)$$

where $\boldsymbol{H}_{(\cdot)}$ denotes the Hessian of the respective potential in the origin. This can be recast as a coupled system of Mathieu-type equations,

$$\ddot{\boldsymbol{x}} + \left(\frac{\Omega_{\mathrm{rf}}}{2}\right)^2 (\mathcal{A} + 2\mathcal{Q}\cos(\Omega_{\mathrm{rf}}t))\,\boldsymbol{x} = 0, \qquad (2.19)$$

where the dimensionless matrices

$$\mathcal{A} = \frac{2Ze}{m\Omega_{\mathrm{rf}}^2}\boldsymbol{H}_{\Phi_{\mathrm{dc}}}, \quad \mathcal{Q} = \frac{Ze}{m\Omega_{\mathrm{rf}}^2}\boldsymbol{H}_{\Phi_{\mathrm{rf}}} \qquad (2.20)$$

are generalisations of the usual scalar Mathieu parameters a and q.

If $\mathcal{A}$ and $\mathcal{Q}$ are simultaneously diagonalisable, the system is easily decoupled by a coordinate transformation $x_i \mapsto \tilde{x}_i$ into the shared eigenbasis, resulting in individual Mathieu equations

$$\ddot{\tilde{x}}_i + \left(\frac{\Omega_{\mathrm{rf}}}{2}\right)^2 (\tilde{a}_i + 2\tilde{q}_i\cos(\Omega_{\mathrm{rf}}t))\,\tilde{x}_i = 0, \quad i = 1, 2, 3, \qquad (2.21)$$

with stability parameters $\tilde{a}_i = 2Ze/(m\,\Omega_{\mathrm{rf}}^2)\,\partial^2\Phi_{\mathrm{dc}}/\partial\tilde{x}_i^2$ and $\tilde{q}_i = Ze/(m\,\Omega_{\mathrm{rf}}^2)\,\partial^2\Phi_{\mathrm{rf}}/\partial\tilde{x}_i^2$. Paul traps for quantum information processing are typically operated in the first stability region, with $|a| \ll |q| \ll 1$.

The decoupled case is widely studied [Lei+03], perhaps in part as both Φ_{dc} and Φ_{rf} are typically well-approximated by electrode-aligned quadrupole potentials in macroscopic linear ("four-rod") Paul traps. In contrast to this, symmetric surface trap designs are usually operated with a static potential Φ_{dc} that has principal axes deliberately not aligned with the trap plane (and as such not aligned with the principal axes[6] of Φ_{rf}) to ensure a non-zero projection of all secular modes of motion onto the cooling laser beams. For simplicity, we still focus on the scalar case here, as we will later combine the effects of DC

---

[6]Note that $\mathcal{A}$ and $\mathcal{Q}$ being simultaneously diagonalisable, such that the equations of motion decouple into scalar Mathieu equations, is a condition on the principal axes of the quickly-varying RF quadrupole potential Φ_{rf}, not those of the resulting pseudopotential that describes the resulting radial confinement. In fact, for typical designs with (approximately) vanishing axial RF curvature, as exemplified by an infinitely long linear trap, the RF-induced pseudopotential is necessarily symmetric as a consequence of Laplace's equation; only the effect of the DC electrode voltages lifts the degeneracy.

and time-averaged RF potentials into a single, effective confining potential, thus obtaining analytic results in what is known as the pseudopotential approximation. If necessary for quantitative calculations, Floquet solutions to the general eq. (2.19) could similarly be found using a technique of infinite-continued matrix inversions [Hou08; Lan+12].

To first order in $\tilde{a}_i$ and $\tilde{q}_i$, the solution to the (classical) equations of motion (2.21) can be approximated as

$$\tilde{x}_i(t) \approx \tilde{X}_i(t) \left(1 + \frac{q_i}{2} \cos(\Omega_{\mathrm{rf}} t) \right), \tag{2.22}$$

where $\tilde{X}_i(t)$ describes the classical trajectory of a harmonic oscillator with some secular mode frequency $\omega_{m,i}$, and the $\cos(\Omega_{\mathrm{rf}} t)$ term describes *micromotion* at the faster time-scale $\Omega_{\mathrm{rf}} \gg \omega_{m,i}$. Assuming the confinement in the radial direction is given entirely by the time-averaged effect of Φ_{rf}, the Mathieu q parameter and secular mode frequency $\omega_{m,i}$ are related by

$$\tilde{q}_i \approx 2\sqrt{2}\, \frac{\omega_{m,i}}{\Omega_{\mathrm{rf}}}. \tag{2.23}$$

This treatment is straightforwardly extended to the quantum case (see e.g. [Lei+03]), though we do not attempt this here. A representative set of parameters for our experiments is $\omega_{m,i} \approx 2\pi \times 3\,\mathrm{MHz}$ and $\Omega_{\mathrm{rf}} \approx 2\pi \times 50\,\mathrm{MHz}$, i.e. $q_i \approx 0.17$.

This *intrinsic* micromotion contribution from the modulation of the ion's secular motion in eq. (2.22) is often neglected, as $q_i \ll 1$ and $\tilde{X}_i$ is small for ions are cooled close to their motional ground state. However, additional electric fields, caused for instance by stray electric charges near the trapping region, can shift the equilibrium position of the harmonic pseudopotential away from the symmetry axis of the RF potential. Such additional, constant offsets to $\tilde{X}_i$ in practice readily exceed the amplitude of the secular motion of laser-cooled ions. The concomitant increase in the micromotion amplitude as described by eq. (2.22) is referred to as *excess micromotion*. In the harmonic region of the trap, the shift $x_{\mathrm{stray},i}$ induced by a stray electric field component $E_{\mathrm{stray},i}$ in direction i can be calculated simply from the restoring force of the harmonic oscillator as

$$x_{\mathrm{stray},i} = \frac{e\, E_{\mathrm{stray},i}}{m\, \omega_{m,i}{}^2} \tag{2.24}$$

where e is the (elementary) charge of the ion. This results in an additional modulation of the ion position by a term of the form $\frac{1}{2} x_{\mathrm{stray},i}\, q_i\, \cos(\Omega_{\mathrm{rf}} t)$. This modulates all ion–laser interactions and has geometric consequences for the collection of spontaneous emission photons (as discussed in chapter 4), and is generally undesired. Chapter 5 introduces a method to detect and subsequently compensate such fields.

2.2.2 *Internal (electronic) state*

The primary mechanism by which a single trapped atomic ion serves as a qubit is by means of its internal electronic state: we simply designate two of the many possible eigenstates of

the electronic Hamiltonian as $|0\rangle$ and $|1\rangle$ (typically, many such assignments are possible; the ones used in this work are described in chapter 3). Typically, we work with species that have only one valence electron, such as singly-ionised earth alkaline metals, as this gives rise to a tidy level structure that is easy to dissipatively manipulate using laser radiation. How this energy level structure comes about is, as far as its role as a qubit substrate is concerned, mostly an irrelevant detail. The Hamiltonian describing the valence electron (mass m_e, charge $-e$, spin $\hbar\boldsymbol{S}$) in the atomic centre-of-mass frame could, for instance, be understood in the central-field approximation as

$$H = \frac{\boldsymbol{p}^2}{2m_e} + V(\boldsymbol{x}) + g_S\mu_B\boldsymbol{S}\cdot\boldsymbol{B}(\boldsymbol{x}), \tag{2.25}$$

where V is the effective radial potential. For our purposes, however, it is sufficient to regard it in its spectral decomposition, that is, as a collection of its eigenstates $|i\rangle$, where i is for now a generic index[7] from some countable set I_A:

$$H_{\mathrm{atom}} = \hbar \sum_{i\in I_A} \omega_i^A |i\rangle\langle i| . \tag{2.26}$$

Here, we have written $(\hbar\omega_i^A)_{i\in I_A}$ for the energies of the respective levels, which we can also take to result from some more complex description than the simple model in eq. (2.25), for instance, involving spin-orbit coupling, relativistic corrections, or just derived from experimental measurements. In the experiment, a static magnetic field, homogeneous across the trap region, is applied to lift the degeneracy in the electronic states through the Zeeman effect. The most natural choice for the state labels I_A is the set of quantum numbers describing each atomic state. Their choice of course depends on whatever approximation is most well-suited to describe the species in question; for $^{88}\mathrm{Sr}^+$ in the context of this thesis, good quantum numbers are the principal quantum number n, the orbital angular momentum L, the total angular momentum J, and its projection onto the magnetic field axis m_J (often written, together with the single-electron orbital angular momentum l and the total spin magnitude $S = 1/2$, in the traditional spectroscopic notation consisting of configuration and term symbol). As we shall see in §2.4.3, transitions between the levels can be excited using resonant electromagnetic radiation, where selectivity can be ensured by an appropriate choice of frequency, polarisation, or a combination of the two.

2.2.3 Quantised motion

The electrical trap potentials are typically designed such that the time-averaged pseudopotential resulting from averaging across the fast dynamics in eq. (2.21) is in good approximation harmonic in all spatial directions. The centre-of-mass motion of a trapped ion

---

[7]Using this description, we have already restricted ourselves to discrete, bound states of the atom – which is not an issue, as we are not concerned with processes such as photoionisation here.

can thus be described as a combination of three quantum-harmonic oscillator modes; one aligned with each of the principal axes of the confining pseudopotential. The ion position operator can be decomposed in terms of the annihilation and creation operators a_i and $a_i^\dagger$ of the individual modes as [Lei+03]

$$\boldsymbol{x} = \sum_i x_{0,i} \left(a_i + a_i^\dagger\right), \quad x_{0,i} = \sqrt{\frac{\hbar}{2m\omega_{m,i}}}, \tag{2.27}$$

where $x_{0,i}$ is the r.m.s. extent (i.e., positional uncertainty) of the ground-state wave function for mass m and secular frequency $\omega_{m,i}$. To take a typical example, for a $^{88}\text{Sr}^+$ ion in a trap with secular frequency $\omega_m = 2\pi \times 3\,\text{MHz}$, $x_0 = 4\,\text{nm}$.

If N ions are confined in the trap in the form of a linear crystal and the Coulomb interaction is linearised around the equilibrium positions, there are $3N$ harmonic oscillator modes; multiple ions participate in each. This is of crucial importance for the use of trapped ions in quantum computing, as the collective motion provides the "quantum bus" across which multi-ion interactions can be engineered[8].

In the context of the experiments with one ion per trap discussed in this thesis, the secular motion, however, only has very peripheral relevance as a source of infidelity in coherent optical operations. Laser cooling (Doppler cooling) leaves all motional modes in good approximation in a thermal state corresponding to a temperature close to the Doppler limit,

$$T_{\text{Doppler}} \approx \frac{\hbar\Gamma}{2\text{k}_B}, \tag{2.28}$$

where Γ is the linewidth of the transition used for cooling. Parametrised by the average occupation number

$$\overline{n}_i = \frac{1}{e^{\frac{\hbar\omega_i}{\text{k}_B T}} - 1}, \tag{2.29}$$

the density matrix for the thermal state of the i-th mode is

$$\rho_{\text{th},i} = \left(1 - e^{-\frac{\hbar\omega_i}{\text{k}_B T}}\right) \sum_{n=0}^{\infty} e^{-\frac{\hbar\omega_i}{\text{k}_B T} n} |n\rangle\langle n| = \frac{1}{1 + \overline{n}_i} \sum_{n=0}^{\infty} \frac{1}{\left(1 + \frac{1}{\overline{n}_i}\right)^n} |n\rangle\langle n|. \tag{2.30}$$

The Doppler limit corresponds to a lower bound on the average thermal occupation number reachable using this simple cooling technique,

$$\overline{n}_{\text{Doppler},i} \approx \frac{1}{e^{2\omega_i/\Gamma} - 1}. \tag{2.31}$$

---

[8]In typical trapping conditions, the spacing between the equilibrium positions of multiple ions is on the order of a few micrometres. The wave-function overlap of the ion is thus negligible, and they can be treated as distinguishable particles without making reference to their quantum statistics. Direct dipole-dipole coupling between the internal state of neighbouring ions in a trap is extremely weak; for instance, a coupling rate of $\sim 1\,\text{mHz}$ between the magnetic dipole moments of two $^{88}\text{Sr}^+$ ions has been observed [Kot+14]. Electric dipole coupling can be engineered to be appreciable by exciting the ions to Rydberg states; techniques for high-fidelity logic gates using this approach are a current topic of research (see e.g. ref. [Zha+20]).

In terms of the coherent harmonic oscillator states $|\alpha\rangle = e^{\alpha a^\dagger - \bar{\alpha} a} |0\rangle = e^{-\frac{|\alpha|^2}{2}} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle$, thermal states can be decomposed as

$$\rho_{\text{th},i} = \frac{1}{\pi \, \bar{n}_i} \int_{\mathbb{C}} e^{-\frac{|\alpha|^2}{\bar{n}_i}} |\alpha\rangle\langle\alpha| \, d^2\alpha. \tag{2.32}$$

2.3 *Photons*

In this context, we regard a photon as a single quantum of excitation in a well-defined electromagnetic field mode. Both quantum mechanics and classical electrodynamics are linear, and the Heisenberg-picture field operators in the quantum field theory of light obey Maxwell's equation. Thus, the propagation of photons through optical systems can be analysed using classical mode solutions [FT20]; in the linear optics case, the quantum nature of light only plays a role when considering the detection statistics (and entanglement with other subsystems, etc.). Concretely, the electric field operator can be expanded into a set of monochromatic modes $u_{k\Pi}(r)$ as [Lou00; SZ97]

$$E(r, t) = i \sum_{k\Pi} \sqrt{\frac{\hbar k c}{2\varepsilon_0 \mathcal{V}}} \, u_{k\Pi}(r) \, e^{-ikct} \, a_{k\Pi} + \text{h. c.}, \tag{2.33}$$

where $a_{k\Pi}$, $a_{k\Pi}^\dagger$ are the annihilation and creation operators obeying, for each mode, the bosonic (harmonic oscillator) commutation relations

$$[a_{k\Pi}, a_{k'\Pi'}] = [a_{k\Pi}^\dagger, a_{k'\Pi'}^\dagger] = 0, \quad [a_{k\Pi}, a_{k'\Pi'}^\dagger] = \delta_{kk'}\delta_{\Pi\Pi'}. \tag{2.34}$$

For plane wave modes, $u_{k\Pi}(r) = \hat{\epsilon}_{\hat{k}\Pi} e^{ik\cdot r}$, and the sum is over modes described by a three-dimensional wavevector k and two polarisations Π, with $\hat{\epsilon}_{\hat{k}\Pi} \perp k$ some choice of basis for the polarisation. Here, we have considered the system to be embedded inside a finite quantisation volume $\mathcal{V}$ as often done in quantum optics to avoid normalisation issues. Later, we will transition to continuum modes ($\mathcal{V} \to \infty$) by making the replacement $\sum_k \cdot \mapsto \sqrt{\frac{\mathcal{V}}{(2\pi)^3}} \int_{\mathbb{R}^3} \cdot \, d^3k$.

Rather than such an expansion across many plane wave modes, a spatio-temporal wavepacket would lend itself more naturally to describing the emission of a single photon from a stationary emitter. It is indeed possible to define the equivalent of a wave function for photons, though this is not without subtleties [SR07; TG66; MSZ05]. Concretely, if $v_{k\Pi}$ are the coefficients of an excitation decomposed into plane waves $u_{k\Pi}(r)$, the wave-packet mode in the continuum limit can be defined as

$$f(r, t) = i\sqrt{\frac{\hbar c}{2\varepsilon_0}} \sum_{\Pi} \int_{\mathbb{R}^3} \sqrt{\frac{k}{(2\pi)^3}} \, v_{k\Pi} \, u_{k\Pi}(r) \, e^{-ikct} d^3k, \tag{2.35}$$

such that if $\{\tilde{a}^\dagger, \tilde{a}\}$ are the creation and annihilation operators describing a photon in this mode, the electric field operator becomes

$$E(r, t) = f(r, t) \, \tilde{a} + \text{h. c.} + \dots, \tag{2.36}$$

where the ellipsis denotes the further orthogonal modes necessary to form a complete basis for the field[9].

2.3.1   *Temporal wave-packet modes*

For quantum information processing, where photons are typically confined in collimated beams, optical fibres, or waveguides, it is useful to focus on the longitudinal frequency/time component along the propagation of the beams, and assume that the transverse parts are in some spatial mode well-behaved enough not to be of any further interest (e.g. a Gaussian beam). Assuming the beam with index n propagates in the z direction and taking the continuum limit for the longitudinal modes, the electric field operator takes the form

$$\boldsymbol{E}_n(z,t) = i \sum_{\Pi} \hat{\boldsymbol{\epsilon}}_{\Pi} \int_0^{\infty} \sqrt{\frac{\hbar k c}{4\pi\varepsilon_0 \mathscr{A}}} \, e^{ik(z-ct)} \, b_{n\Pi}(k) \, dk + h.\,c., \tag{2.37}$$

where $\mathscr{A}$ is the formal quantisation mode area for the transversal modes[10], and we have written $b_{n\Pi}^{\dagger}(k)$ for the creation operator for beam mode n, polarisation Π, and angular wavenumber k for visual contrast with the plane-wave free-space operators $a_{k\Pi}^{\dagger}$. The canonical commutation relations are

$$\begin{aligned}
\left[b_{n\Pi}(k), b_{n'\Pi'}(k')\right] &= \left[b_{n\Pi}^{\dagger}(k), b_{n'\Pi'}^{\dagger}(k')\right] = 0, \\
\left[b_{n\Pi}(k), b_{n'\Pi'}^{\dagger}(k')\right] &= \delta_{nn'} \, \delta_{\Pi\Pi'} \, \delta(k-k').
\end{aligned} \tag{2.38}$$

A one-photon wavepacket can be created by a linear superposition

$$B_{n\Pi}^{\dagger} := \int_0^{\infty} \tilde{\beta}(k) \, b_{n\Pi}^{\dagger}(k) \, dk, \tag{2.39}$$

of these modes specified by an envelope function $\tilde{\beta}(k) \in \mathbb{C}$ with $\left\|\tilde{\beta}\right\|^2 = \int_0^{\infty} |\tilde{\beta}(k)|^2 dk = 1$. Assuming that the radiation is narrow-band, we can extend the lower integration bound to $-\infty$ and directly introduce temporal modes given by the Fourier-transformed envelope function

$$\beta(t) = \frac{1}{\sqrt{2\pi c}} \int_{-\infty}^{\infty} \tilde{\beta}(k) e^{-ikct} dk \tag{2.40}$$

---

[9]A subtle point here is that the mode functions $\boldsymbol{f}(\boldsymbol{r},t)$ are non-orthogonal even for a unitary change of basis via the coefficients $v_{k\Pi}$, due to the $\sqrt{k}$ factor in eq. (2.35). This can be remedied through a redefinition of the inner product in terms of appropriate dual modes [SR07], though in the present case we are concerned only with narrow-band radiation where $\sqrt{k}$ is approximately constant.

[10]In practice, the transversal modes are also unconfined free-space modes, and we could similarly take the continuum limit – as we are only ever concerned with one transversal mode in the abstract here, however, this would just be an unnecessary complication.
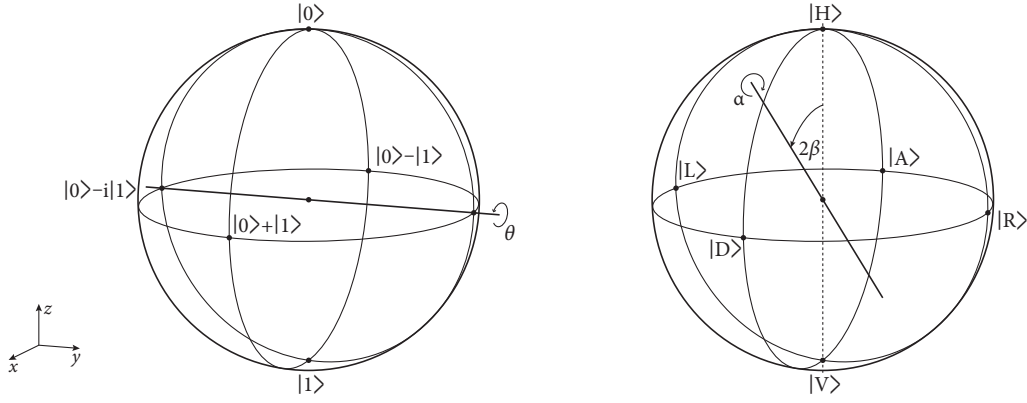
Figure 2.3: Two of many ways to encode a qubit in a photonic mode. *Left:* Encoding in two spatial modes $|0\rangle$ and $|1\rangle$. In our convention, a beamsplitter with reflectance $(\sin\theta)^2$ rotates the Bloch sphere around the y-axis. *Right:* Encoding in the polarisation of a photon. A wave plate of retardance α with fast axis at angle β to the horizontal direction rotates the Bloch sphere as shown. An alternative convention is frequently used which assigns circular polarisation to the poles, matching the Poincaré sphere from classical optics.

such that the creation of the above wavepacket can equivalently be written as

$$
B_{n\Pi}^{\dagger} = \int\limits_{-\infty}^{\infty} \beta(t)\, b_{n\Pi}^{\dagger}(t)\, \mathrm{d}t \tag{2.41}
$$

with the Fourier-conjugate time-domain creation and annihilation operators $b_{n\Pi}^{\dagger}(t)$, $b_{n\Pi}(t)$, for which similarly

$$
\begin{aligned}
\left[b_{n\Pi}(t), b_{n'\Pi'}(t') \right] &= \left[b_{n\Pi}^{\dagger}(t), b_{n'\Pi'}^{\dagger}(t') \right] = 0, \\
\left[b_{n\Pi}(t), b_{n'\Pi'}^{\dagger}(t') \right] &= \delta_{nn'}\, \delta_{\Pi\Pi'}\, \delta(t - t').
\end{aligned} \tag{2.42}
$$

Such temporal wave-packet modes are particularly useful in the context of analysing the photonic Bell-basis measurement apparatus here (§4.4) not only because they naturally capture the case of a propagating beam where there is no interaction across different times t, but also because the measurement operator for time-resolved detection of a single photon in mode $n\Pi$ at time t is simply given by $b_{n\Pi}(t)$.

2.3.2 *Photonic qubits*

By picking two field states, usually an excitation of the field in orthogonal modes, and restricting our attention to states spanned by them, we can define an effective two-level system, or qubit. Many such qubit encodings are possible, and indeed used in practice. One important consideration is the ability to manipulate the two states, e.g. to apply unitary rotations, or to perform projective measurements in a particular basis. Several possible qubit encodings are illustrated in fig. 2.3.

An obviously orthogonal choice for two qubit basis states are excitations in two different spatial modes that have no overlap, such as two well-separated collimated beams; this

is sometimes also referred to as dual-rail encoding. To implement x–y-plane rotations, the beam modes corresponding to the qubit basis need to be brought together at a beam-splitter, where they interfere; z-axis rotations can be implemented by differential phase shifts.

In the laboratory, we usually think of beamsplitters as being parametrised by their re-flectance R and transmittance T ($R + T = 1$, e.g. $R = T = \frac{1}{2}$ for the ideal "half-silvered mirror"). Without further information, it is not clear what the phase of the output relative to the input beams should be (apart from constraints due to unitarity). For our narrow-band use case the difference is trivial[11], and in either case would not be relevant, as the optical part of our apparatus only consists of a diverging network of beamsplitters and we do not have any absolute phase reference for the output beams. Here, we choose a phase convention such that the unitary matrix for a beamsplitter is

$$U_{BS}(\theta) = \begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix}, \tag{2.43}$$

where $(\cos\theta)^2 = T$ and $(\sin\theta)^2 = R$. This is a rotation around the y axis of the Bloch sphere.

One particularly straightforward choice is to encode the qubit in the polarisation of the photon, which is naturally a two-level system. Unitary rotations are easy to implement by inserting birefringent materials into the beam path (e.g. wave plate retarders in a rotation mount transversal to the beam). Most photodetectors are (approximately) polarisation-insensitive, but polarisation-basis measurements are easily realised by placing the detector behind e.g. a polarising beamsplitter cube (PBS, which effectively converts between polar-isation and spatial mode[12] encodings).

Here, we define the linear polarisation states $|H\rangle$ and $|V\rangle$ as the computational basis states on the z axis of the Bloch sphere, as it is natural in the context of beams manipu-lated using PBS cubes (though it differs from the classical optics convention of drawing the Poincaré sphere with circular polarisation on the north and south poles). The unitary action of a general birefringent element with retardance α, the fast axis of which is at angle β to the horizontal direction, is

$$U_{\mathrm{wp}}(\alpha, \beta) = \begin{pmatrix} (\cos\beta)^2 + e^{-i\alpha}(\sin\beta)^2 & (1 - e^{-i\alpha})\cos\beta\sin\beta \\ (1 - e^{-i\alpha})\cos\beta\sin\beta & e^{-i\alpha}(\cos\beta)^2 + (\sin\beta)^2 \end{pmatrix}; \tag{2.44}$$

the common half- and quarter-wave plates correspond to $\alpha = \pi$ and $\alpha = \frac{\pi}{2}$, respectively.

---

[11] In more general cases, such as a white-light interferometer, the phase imparted by the beamsplitter is not necessarily irrelevant; see ref. [Hamoo] for a summary.

[12] One could argue that the two output modes of a PBS cube do not strictly correspond to dual-rail encod-ing as the modes have orthogonal polarisations. If desired, however, the polarisations can easily be equalised by inserting a half-wave plate into one of the output arms.

While polarisation encoding makes photonic qubits easy to manipulate, such states are also sensitive to unintended birefringence in the optical path. For most optical elements, residual birefringence will be relatively stationary and thus easily corrected. Optical fibres are a notable exception, however; they exhibit considerable changes in stress-induced birefringence when subjected to temperature and humidity changes or to mechanical disturbances. Polarisation-maintaining fibres, where the fibre core is under considerable anisotropic stress by design, suppress coupling of linear input polarisation between a pair of orthogonal design axes (e.g. horizontal and vertical). Empirically, however, their action on other input states is considerably more sensitive to disturbances (causing for instance drifts between diagonal and circular polarisation). To transmit polarisation-encoded qubits in fibre, regular single-mode fibre is thus preferred, augmented by active polarisation stabilisation as necessary to cover longer distances (which is routinely implemented across many kilometres of deployed fibre in the context of quantum key distribution in metropolitan-area networks, see e.g. ref. [Tre+09]).

The photon frequency is yet another possible qubit encoding, which has also been used in context of atom-photon entanglement – if the frequency difference between two wavepacket modes is much larger than their bandwidth, they are approximately orthogonal. Frequency encoding has the potential advantage that a splitting on the ~ GHz scale is very small compared to ~ 500 THz optical frequencies, so most interactions with propagation media, etc., will be very similar for the two states, leading to robust transmission. However, this also means that frequency-encoded qubits are hard to manipulate efficiently, requiring narrow-band filters (e.g. based on optical cavities) for the analysis of the photon state or efficient Bell-basis measurements.

Many other photonic qubit/qudit encodings exist beyond those mentioned here, for instance in different time bins, in the transversal beam modes, including orbital angular momentum beams [Mai+01], in different temporal photon envelopes [Bre+15], etc. For ion–photon entanglement, it will also be quite common for photons to be entangled with the atom in more than one degree of freedom, for instance both in polarisation (due to selection rules) and frequency (e.g. due to differing Zeeman shifts). As long as there is at least one modality in which the states are orthogonal, they are also orthogonal overall, so this appearance of multiple degrees of freedom is benign.

Generally, photonic systems will often exhibit non-negligible transmission loss, and the efficiency of single-photon detectors is limited. This can be straightforwardly expressed in the full second-quantisation picture as an admixture of the vacuum state. In experiments post-selected on the presence of a given number of photons, e.g. when discussing ion–photon entanglement, it is often preferable to stay in the encoded-qubit picture, however, and track the overall probability of observing a photon in that subspace separately or through the use of sub-normalised density operators.

2.4 Atom-field interactions

We now consider the interaction that underpins all the techniques related to the internal electronic state of trapped ions used in this work, that between the internal state and an external electromagnetic field. This encompasses both coherent manipulation, incoherent scattering, and the generation of ion–photon entanglement through spontaneous emission. We will first treat the generic Hamiltonian of the system in a fully quantum-mechanical fashion to elucidate the different modes of interaction, and only then discuss the special cases of interest here.

2.4.1 Interaction of a bound electron with the multipolar e.m. field

We are interested in the non-relativistic Hamiltonian that describes a single electron (mass m_e, charge $-e$, spin $\hbar S$), the electromagnetic field, and their interaction. The precise form of this interaction has historically been a remarkably persistent source of disagreement and confusion (see e.g. refs. [Lam52, sec. 60], [AM84], [RF17], [And+18], [FLM19]). Here, we start from[13] the canonical Hamiltonian obtained by transforming the minimal-coupling Lagrangian for the electron, the field and the Lorentz force,

$$H = \frac{(\boldsymbol{p} + e\boldsymbol{A}(\boldsymbol{x}))^2}{2m_e} + V(\boldsymbol{x}) + \frac{1}{2}\int_{\mathbb{R}^3}\left(\varepsilon_0 \boldsymbol{E}_\perp(\boldsymbol{r})^2 + \frac{1}{\mu_0}\boldsymbol{B}(\boldsymbol{r})^2\right)\mathrm{d}^3\boldsymbol{r} + g_S\mu_B \boldsymbol{S}\cdot\boldsymbol{B}(\boldsymbol{x}), \quad (2.45)$$

where $\boldsymbol{x}$ is the electron position, $\boldsymbol{p}$ its canonical momentum ($\boldsymbol{p} = \boldsymbol{p}_{\mathrm{kin}} - e\boldsymbol{A}(\boldsymbol{x})$, $\boldsymbol{p}_{\mathrm{kin}} := m_e\dot{\boldsymbol{r}}$), and we have added an additional term describing the interaction between the electron spin and magnetic field as quantified by the Bohr magneton μ_B and the electron g-factor $g_S \approx 2$. $\boldsymbol{E}_\perp$ is the transversal (divergence-, i.e. source-free) part of the electric field[14]; we have eliminated the longitudinal (curl-free) part by choosing to work in the Coulomb gauge ($\boldsymbol{\nabla} \cdot \boldsymbol{A} = 0$) and writing the interaction with the rest of the atom as an effective potential $V(\boldsymbol{x})$. This is a rather simplified model of the interactions internal to the atom, but should do just fine for purposes of exposition. For now, we have also assumed the nucleus to be infinitely heavy and in the origin.

Equation (2.45) is sometimes referred to as the "$\boldsymbol{p}\cdot\boldsymbol{A}$ Hamiltonian", after the expansion $(\boldsymbol{p} + e\boldsymbol{A}(\boldsymbol{x}))^2 = \boldsymbol{p}^2 + 2e\boldsymbol{p}\cdot\boldsymbol{A}(\boldsymbol{x}) + e^2\boldsymbol{A}(\boldsymbol{x})^2$ (in the Coulomb gauge, $[\boldsymbol{p}, \boldsymbol{A}] = 0$). The key ingredient in its analysis is what has become known as the Power–Zienau–Woolley (PZW) transformation [PZ59; Woo71], a unitary change of frame for Hamiltonian and states that will allow us to shift the electron-field interaction from the $\boldsymbol{p} \cdot \boldsymbol{A}$ term evaluated at the

---

[13]The treatment here mostly follows ref. [BPT74] and [Lou00, ch. 4], though I have chosen to add the electron spin term for clarity, as magnetic dipole transitions between Zeeman states are ubiquitous in trapped ions.

[14]As in the rest of this thesis, no typographic distinction is drawn between quantities that are operators and those that are not, and in the later semi-classical approximation, various quantities such as the electric and magnetic fields will assume an explicit time dependence.

electron position to under a field integral, such that we can conveniently take its multipole expansion (see e.g. refs. [And+20; Woo20] for a modern perspective). To start with, we define a quantity

$$P(r) := -ex \int_0^1 \delta(r - sx)\, ds, \qquad (2.46)$$

which could be interpreted as a polarisation density (or, at least, its longitudinal component [And+18]), as $\nabla \cdot P(r) = -\rho(r)$ for the charge density ρ. Here, however, we can be content with its role as part of a generator S

$$S := \int_{\mathbb{R}^3} P(r) \cdot A(r)\, d^3r, \qquad (2.47)$$

which we use to define the unitary transformation $H \mapsto H' := e^{\frac{i}{\hbar}S} H e^{-\frac{i}{\hbar}S}$ (such that, for states, $|\psi\rangle \mapsto e^{\frac{i}{\hbar}S}|\psi\rangle$). By expanding the conjugation by exponentials as a series of commutators, we obtain the result, the multipolar or PZW Hamiltonian, as

$$H' = \underbrace{\frac{p^2}{2m_e} + V(x)}_{=:H_{\text{atom}}} + \underbrace{\frac{1}{2}\int_{\mathbb{R}^3}\left(\varepsilon_0 E_\perp(r)^2 + \frac{1}{\mu_0}B(r)^2\right)d^3r}_{=:H_{\text{field}}} + H_{\text{cpl}}, \qquad (2.48)$$

where we can recognise the usual terms of the free atom and field, H_{atom} and H_{field}, and the new terms H_{cpl} are [BPT74]

$$H_{\text{cpl}} := \underbrace{-\int_{\mathbb{R}^3} P_\perp(r) \cdot E_\perp(r)\, d^3r + \frac{1}{2\varepsilon_0}\int_{\mathbb{R}^3} P_\perp(r)^2\, d^3r}_{=:H_{\text{cpl,el}}}$$

$$\underbrace{-\int_{\mathbb{R}^3} M(r) \cdot B(r)\, d^3r + \frac{e^2}{2m_e}\left(x \times \int_0^1 sB(sx)ds\right)^2 + g_S\mu_B S \cdot B(x)}_{=:H_{\text{cpl,mag}}}, \qquad (2.49)$$

where we have also introduced a formal magnetisation density[15] analogous to $P(r)$,

$$M(r) := -\frac{e}{m_e}(x \times p)\int_0^1 s\delta(r - sx)\, ds. \qquad (2.50)$$

The $H_{\text{cpl,el}}$ part describes electric effects; the first term being the atom-field interaction we are seeking, the second a field-independent correction to the atomic self-energy, which

---

[15]There are some subtleties here regarding the ordering of operators, as well as the fact that the canonical momentum p is not quite the usual kinetic momentum $m_e\dot{r}$ – which we shall unceremoniously ignore here, as magnetic interactions are not the focus of this work. If the interaction of ions with structured magnetic fields is of interest, these issues should be critically reviewed (see e.g. ref. [CDG97, chapter IV.C, p. 280ff.] for a discussion of this, where the transformations are applied to the Lagrangian before canonical quantisation instead).

contributes to the Lamb shift [Mil76], but which we will merge into the atomic part here. The $H_{\text{cpl,mag}}$ part describes the magnetic effects, where the second summand is a generalised diamagnetic energy, quadratic in $\boldsymbol{B}$.

2.4.2   Multipole expansion

Armed with eq. (2.49), which makes the dependence on the gauge-invariant transversal fields $\boldsymbol{E}_\perp$ and $\boldsymbol{B}$ explicit, we can now consider the atom-field interaction in multipole expansion. Consider the electric term

$$
\begin{aligned}
-\int_{\mathbb{R}^3} \boldsymbol{P}_\perp(\boldsymbol{r}) \cdot \boldsymbol{E}_\perp(\boldsymbol{r})\, \mathrm{d}^3 \boldsymbol{r} &= e\boldsymbol{x} \cdot \int_0^1 \boldsymbol{E}_\perp(s\boldsymbol{x})\, \mathrm{d}s \\
&= e\boldsymbol{x} \cdot \int_0^1 \left(\boldsymbol{E}_\perp(0) + s\nabla \boldsymbol{E}_\perp(0)\,\boldsymbol{x} + \frac{s^2}{2!} \ldots \right) \mathrm{d}s \qquad (2.51) \\
&= e\boldsymbol{x} \cdot \left(\boldsymbol{E}_\perp(0) + \frac{1}{2!}\nabla \boldsymbol{E}_\perp(0)\,\boldsymbol{x} + \frac{1}{3!} \ldots \right),
\end{aligned}
$$

where the dots elide the higher-order terms in the Taylor expansion[16] of $\boldsymbol{E}_\perp(\boldsymbol{r})$ at $\boldsymbol{r} = 0$. The lowest-order term gives the electric dipole (E1) interaction Hamiltonian,

$$
H_{\text{E1}} = e\boldsymbol{x} \cdot \boldsymbol{E}_\perp(0), \qquad (2.52)
$$

and the second-order term is the electric quadrupole (E2) interaction,

$$
H_{\text{E2}} = \frac{1}{2} e\boldsymbol{x} \cdot \nabla \boldsymbol{E}_\perp(\boldsymbol{0})\boldsymbol{x}. \qquad (2.53)
$$

We can similarly expand $\boldsymbol{B}(\boldsymbol{r})$ around $\boldsymbol{r} = 0$ in the magnetic term $H_{\text{cpl,mag}}$, and obtain, to lowest order in the electron position, the familiar magnetic dipole (M1) interaction

$$
H_{\text{M1}} = \left(\underbrace{\frac{e}{2m_e}(\boldsymbol{x} \times \boldsymbol{p})}_{=\mu_B \boldsymbol{L}} + g_S \mu_B \boldsymbol{S} \right) \cdot \boldsymbol{B}(0), \qquad (2.54)
$$

in addition to a diamagnetic shift $\frac{e^2}{8m_e}\left(\boldsymbol{x} \times \boldsymbol{B}(0)\right)^2$, which is quadratic in $\boldsymbol{B}$, and negligible for relevant field strengths.

These low-order truncations are in general good approximations for optical transitions[17] (and excellent for radio-frequency magnetic transitions), as each additional order in the expansion is scaled by an extra factor on the order of $r/\lambda \ll 1$, where λ is the wavelength, and r a typical value for $|\boldsymbol{x}|$. Considering the coupling to the same radiation field

---

[16]The contraction with higher-rank tensors of partial derivatives occurring in these terms means they are more conveniently written in component form.

[17]Next-order electric octupole (E3) transitions are sometimes interrogated in trapped-ion optical atomic clocks, for instance in $^{171}\text{Yb}^+$[Hun+12; Kin+12].

($|\boldsymbol{B}| = |\boldsymbol{E}|/c$), the M1 coupling is typically much weaker than the E1 coupling due to the extra $p/(m_e c)$ factor (where p is a typical value for $|\boldsymbol{p}|$), such that we can often – certainly in the scope of this thesis – consider optical transitions as mediated by the electric field only.

2.4.3   Coupling between atomic levels

Note that the Hamiltonian describing the interaction (H_{cpl}), commutes neither with the atomic (H_{atom}) nor field (H_{field}) parts – none of the low-order multipole terms H_{E1}, H_{E2}, H_{M1} do. In the quantum picture, we would of course expect no less, as e.g. spontaneous emission couples even to the vacuum field. Thankfully, in our case, the radiation fields are typically small, so the interactions can be successfully described as acting on the bare atomic states using perturbation theory (with the high-field regime, where this approach breaks down, typically starting around $\sim 10^{17}$ W/m^2 [And+20]).

Writing the undisturbed atomic Hamiltonian in its spectral decomposition (eq. (2.26)), and inserting the decomposition of the identity[18] $\mathbb{1} = \sum_{i \in I_A} |i\rangle\langle i|$ on either side, we can write, say, the dipole coupling Hamiltonian as

$$H_{\text{E1}} = \sum_{i,j \in I_A} e\,\langle i|\boldsymbol{x}|j\rangle \cdot \boldsymbol{E}_\perp(0)\,|i\rangle\langle j|\,. \tag{2.55}$$

This is a useful sleight, as we can now inject any knowledge about the structure of the eigenstates $|i\rangle$ and $|j\rangle$. In particular, in the atomic species of interest here the spin-orbit interaction leads to an atomic structure captured well by the *LS* coupling model, where the eigenvalue J of the total angular momentum operator $\boldsymbol{J}^2$ and m_J of its z-projection J_z are "good" quantum numbers. In such a case, where the states can be labelled $|\alpha J m_J\rangle$ (where α represents other, non-angular-momentum quantum numbers), we can bring the Wigner–Eckert theorem to bear [BS68]. For this, we expand $\boldsymbol{x}$ in terms of the irreducible spherical tensor operators of rank 1, $\{T_q^{(1)}\}_{q=-1,0,1}$, such that $\boldsymbol{x} = \sum_{q=-1}^1 T_q^{(1)} \hat{\boldsymbol{e}}_q$, where $\{\hat{\boldsymbol{e}}_q\}_q$ are the spherical basis vectors

$$\begin{aligned} \hat{\boldsymbol{e}}_1 &= -(\hat{\boldsymbol{x}} + i\hat{\boldsymbol{y}})/\sqrt{2}, \\ \hat{\boldsymbol{e}}_0 &= \hat{\boldsymbol{z}}, \\ \hat{\boldsymbol{e}}_{-1} &= (\hat{\boldsymbol{x}} - i\hat{\boldsymbol{y}})/\sqrt{2}. \end{aligned} \tag{2.56}$$

---

[18] Implicit throughout is that atomic projectors $|i\rangle\langle i|$ act as the identity on the photonic part of the overall tensor product Hilbert space $H = H_{\text{atom}} \otimes H_{\text{photon}}$.

We can then evaluate the matrix element from eq. (2.55) for two states $|i\rangle = |\alpha_i J_i m_{Ji}\rangle$, $|j\rangle = |\alpha_j J_j m_{Jj}\rangle$ as[19]

$$
\begin{aligned}
\langle i|\boldsymbol{x}|j\rangle &= \sum_{q=-1}^{1} \langle \alpha_i J_i m_{Ji}|T_q^{(1)}|\alpha_j J_j m_{Jj}\rangle \, \hat{\boldsymbol{e}}_q \\
&= \langle \alpha_i J_i\|T^{(1)}\|\alpha_j J_j\rangle \sum_{q=-1}^{1} \langle J_i, m_{Ji}, 1, q|J_j, m_{Jj}\rangle \, \hat{\boldsymbol{e}}_q \\
&= -(-1)^{J_i-m_{Jj}} \sqrt{2J_j+1} \, \langle \alpha_i J_i\|T^{(1)}\|\alpha_j J_j\rangle \sum_{q=-1}^{1} \begin{pmatrix} J_i & 1 & J_j \\ m_{Ji} & q & -m_{Jj} \end{pmatrix} \hat{\boldsymbol{e}}_q
\end{aligned}
\tag{2.57}
$$

where $\langle \alpha_i J_i\|T^{(1)}\|\alpha_j J_j\rangle$ denotes the reduced matrix element, $\langle J_i, m_{Ji}, 1, q|J_j, m_{Jj}\rangle$ the Clebsch–Gordan coefficient, and the 2×3 terms in parentheses the Wigner 3–j symbol[20]. The latter produces the dipole selection rules: The coupling is non-zero only if $\Delta m_J := m_{Jj} - m_{Ji} \in \{-1, 0, 1\}$, corresponding to σ^+ ($q = 1$), π ($q = 0$) and σ^- ($q = -1$) transitions.

2.4.4 *Classical radiation: coherent gates, scattering*

Consider now the interaction of two atomic levels, which we will identify as $|0\rangle$ and $|1\rangle$, with a monochromatic classical radiation field $\boldsymbol{E}_\perp(0) = E_0 \cos(\omega_l t + \phi) \, \hat{\boldsymbol{\epsilon}}$, for instance as produced by a laser. In the interaction picture w.r.t. the atomic levels, the relevant term in H_{E1} from (2.55) becomes

$$
H_I = -\hbar \frac{\Omega}{2} |1\rangle\langle 0| \, e^{i(\Delta t + \phi)} + \text{h. c.},
\tag{2.58}
$$

where we have introduce the Rabi frequency $\Omega := -\frac{e E_0}{\hbar} \langle 1|\boldsymbol{x}|0\rangle \cdot \hat{\boldsymbol{\epsilon}}$ and the detuning $\Delta := \omega_{10} - \omega_l$ as the difference between atomic transition and laser frequencies. On resonance ($\Delta = 0$), the interaction drives rotations around an axis in the xy plane of the Bloch sphere, generating after time t the unitary evolution

$$
R_{xy}(\vartheta, \phi) = \exp\left(-i\frac{\vartheta}{2}\left((\cos\phi)X + (\sin\phi)Y\right)\right),
\tag{2.59}
$$

where $\vartheta = \Omega t$. All elementary quantum gates used in this work are of this type, though mostly through the equivalent interaction mediated by H_{E2}, where $\Omega := -\frac{e E_0}{2\hbar} \langle 1|(\boldsymbol{x} \cdot \boldsymbol{k})(\hat{\boldsymbol{\epsilon}} \cdot \boldsymbol{x})|0\rangle$ now explicitly depends on the wavevector $\boldsymbol{k}$ of what we have assumed is a plane-wave classical field.

---

[19] Recall that in the convention use throughout this thesis, hats indicate unit vectors, not quantum operators!

[20] Various sign- and normalisation conventions exist here; I follow that of ref. [BS68].

By treating the electric field as fully classical, we have assumed that the natural linewidth Γ of the excited state $|1\rangle$, as set by spontaneous emission (eq. (2.74)), is negligible compared to the Rabi frequency Ω. For weak excitation, where Γ dominates, the ion scatters photons at a rate $R = \Gamma \rho_{11}(\Delta)$, and the steady-state excited state population ρ_{11} follows a Lorentzian frequency response

$$\rho_{11}(\Delta) = \frac{s/2}{1 + s + \left(\frac{2\Delta}{\Gamma}\right)^2}, \tag{2.60}$$

broadened according to the saturation parameter $s := 2\,\Omega^2/\Gamma^2$ [Lei+03]. This is the interaction used for ion state detection through resonance fluorescence and, after taking the frequency shift due to the ion centre-of-mass motion into account, for Doppler cooling.

2.4.5   Spontaneous emission

We return to a fully quantised treatment of the electric field. Even though spontaneous emission is a ubiquitous phenomenon, its proper, full treatment is surprisingly complex. This complexity is in large part due to the involvement of a continuum of vacuum modes, with the usual associated challenges such as the renormalisation of state energies. Here, the emission bandwidth is very narrow compared to the transition frequency (the state lifetime is very long compared to an optical period), which allows us to obtain an analytic approximation following what has become the standard treatment in the tradition of a seminal paper by Weisskopf and Wigner [WW27], as for example presented in ref. [SZ97, ch. 6.3, 10.A], although it should be noted that less heuristic methods, e.g. using a resolvent formalism [CDG97, p. 189ff.], are also possible. We furthermore carry over the assumption of an immobile (infinitely heavy) nucleus in the coordinate origin from §2.4.

We consider the general case of an atom in an excited state $|e\rangle$ and a number of possible target states $|i\rangle \in I_L$ to which it can decay under spontaneous emission of a photon. In particular, we will assume the states to be connected through an electrical dipole interaction given by the Hamiltonian (2.55).[21] From the form of the coupling term for each transition, that of one atomic dipole in a fixed orientation coupling to the field in the origin, it should not come as a surprise that the pattern of radiation observed during spontaneous decay is the same as for an equivalent classical dipole (see fig. 4.2). One can make this analogy to a classical oscillating dipole more explicit by transforming the system Hamiltonian in an appropriate fashion and considering the generic propagator for the e.m. vector potential (cf. [KM76], [Mil94, p. 125ff.]). Here, however, we will just consider the interaction with a generic electric radiation field $E_\perp$, from which the correct angular dependence will emerge naturally.

---

[21] This treatment could be applied just the same to higher-order multipole interactions; the results would be identical in everything but the spatial dependence of the emission patterns.

We expand the electric field operator into a number of plane-wave modes,

$$E_\perp(r) = i\sum_{k\Pi} \sqrt{\frac{\hbar kc}{2\varepsilon_0 \mathscr{V}}} \hat{\epsilon}_{\hat{k}\Pi} e^{ik\cdot r} a_{k\Pi} + \text{h. c.,} \qquad (2.61)$$

where the sum is over modes described by a three-dimensional wavevector k, with $k := |k|$, $k = k\hat{k}$, and one of two polarisations Π, the basis vectors $\hat{\epsilon}_{\hat{k}\Pi} \perp k$ we will leave generic for now, and $\mathscr{V}$ is a large quantisation volume (we will presently consider the limit $\mathscr{V} \to \infty$). In the rotating wave approximation, the dipole interaction Hamiltonian from eq. (2.55) becomes

$$H_{\text{int}} = i\hbar \sum_{i\in I_L} \sum_{k\Pi} \tilde{g}_k^{(\mathscr{V})} \left(g_{i,\hat{k}\Pi} |i\rangle\langle e| a_{k\Pi}^\dagger - \text{h. c.}\right), \qquad (2.62)$$

where we have introduced a shorthand for the field strength prefactor corresponding to the mode,

$$\tilde{g}_k^{(\mathscr{V})} := \sqrt{\frac{kc}{2\hbar\varepsilon_0 \mathscr{V}}}, \qquad (2.63)$$

and the coupling strengths $g_{i,\hat{k}\Pi}$ arise from the dipole matrix element $d_{ie} := -e\langle i|x|e\rangle = -e\langle \alpha_i J_i \| T^{(1)} \| \alpha_e J_e\rangle C_i \hat{e}_{q_i}$ from eq. (2.57) (where $\hat{e}_{q_i}$ is the one spherical basis vector for which the Clebsch–Gordan coefficient $C_i := \langle J_i, m_{Ji}, 1, q_i | J_e, m_{Je}\rangle$ does not vanish) as

$$g_{i,\hat{k}\Pi} := d_{ie} \cdot \overline{\hat{\epsilon}_{\hat{k}\Pi}} = \underbrace{d_{ie}}_{\substack{\text{scalar dipole} \\ \text{strength}}} \underbrace{\hat{d}_{ie} \cdot \overline{\hat{\epsilon}_{\hat{k}\Pi}}}_{\substack{\text{geometric} \\ \text{part}}}. \qquad (2.64)$$

In the Schrödinger picture, the initial state $|\psi(t=0)\rangle = |e\rangle \otimes |0\rangle$, where $|0\rangle$ is the photonic vacuum state, evolves according to $\frac{d}{dt}|\psi(t)\rangle = -\frac{i}{\hbar}H|\psi(t)\rangle$, $H = H_{\text{atom}} + H_{\text{field}} + H_{\text{int}}$. Writing $\tilde{c}_e(t)$, $\tilde{c}_{i,k\Pi}(t) \in \mathbb{C}$ for the coefficients of the interaction-picture state

$$|\tilde{\psi}(t)\rangle := e^{\frac{i}{\hbar}(H_{\text{atom}}+H_{\text{field}})t}|\psi(t)\rangle = \tilde{c}_e(t)|e\rangle \otimes |0\rangle + \sum_{i\in I_L}\sum_{k\Pi} \tilde{c}_{i,k\Pi}(t)|i\rangle \otimes a_{k\Pi}^\dagger|0\rangle, \qquad (2.65)$$

we have

$$|\psi(t)\rangle = \tilde{c}_e(t) e^{-i\omega_e t}|e\rangle \otimes |0\rangle + \sum_{i\in I_L}\sum_{k\Pi} \tilde{c}_{i,k\Pi}(t) e^{-i(\omega_i+kc)t}|i\rangle \otimes a_{k\Pi}^\dagger|0\rangle, \qquad (2.66)$$

and can describe the time evolution of the system in terms of the slowly varying coefficients

$$\frac{d}{dt}\tilde{c}_e(t) = -\sum_{i\in I_L}\sum_{k\Pi} \tilde{g}_k^{(\mathscr{V})} \overline{g_{i,\hat{k}\Pi}} e^{i(\omega_{ei}-kc)t} \tilde{c}_{i,k\Pi}(t), \qquad (2.67)$$

$$\frac{d}{dt}\tilde{c}_{i,k\Pi}(t) = \tilde{g}_k^{(\mathscr{V})} g_{i,\hat{k}\Pi} e^{-i(\omega_{ei}-kc)t} \tilde{c}_e(t), \qquad (2.68)$$

where we have introduced the difference frequency of the atomic eigenstates $\omega_{ei} := \omega_e - \omega_i$. From the initial conditions $\tilde{c}_e(0) = 1$ and $\tilde{c}_{i,k\Pi}(0) = 0$ ($\forall i, k, \Pi$), we can formally solve

eq. (2.68) by integration and insert it into eq. (2.67), giving a single integro-differential equation for the excited state amplitude $\tilde{c}_e(t)$:

$$\frac{\mathrm{d}}{\mathrm{d}t}\tilde{c}_e(t) = -\sum_{i\in I_L}\sum_{k\Pi}(\tilde{g}_k^{(\mathcal{V})})^2|g_{i,\hat{k}\Pi}|^2\int_0^t \tilde{c}_e(t')e^{i(\omega_{ei}-kc)(t-t')}\mathrm{d}t'. \tag{2.69}$$

Transitioning to continuum modes we have, with $\tilde{g}_k := \sqrt{\frac{\mathcal{V}}{(2\pi)^3}}\,\tilde{g}_k^{(\mathcal{V})} = \sqrt{\frac{kc}{2(2\pi)^3\hbar\varepsilon_0}}$,

$$\frac{\mathrm{d}}{\mathrm{d}t}\tilde{c}_e(t) = -\sum_{i\in I_L}\sum_{\Pi}\int_{\mathbb{R}^3} \tilde{g}_k^2|g_{i,\hat{k}\Pi}|^2\int_0^t \tilde{c}_e(t')e^{i(\omega_{ei}-kc)(t-t')}\mathrm{d}t'\,\mathrm{d}^3\boldsymbol{k}. \tag{2.70}$$

As the terms that depend on the magnitude k and those that depend on the geometric aspects $(\hat{\boldsymbol{k}}, \Pi)$ neatly separate, it is convenient to employ spherical coordinates:

$$\frac{\mathrm{d}}{\mathrm{d}t}\tilde{c}_e(t) = -\sum_{i\in I_L}\left(\sum_{\Pi}\int_{\mathcal{S}^2}|g_{i,\hat{k}\Pi}|^2\mathrm{d}^2\hat{\boldsymbol{k}}\right)\int_0^\infty k^2\tilde{g}_k^2\int_0^t \tilde{c}_e(t')e^{i(\omega_{ei}-kc)(t-t')}\mathrm{d}t'\,\mathrm{d}^3 k. \tag{2.71}$$

Independent of the concrete choice of polarisation basis, the integral over unit sphere surface elements $\mathrm{d}^2\hat{\boldsymbol{k}}$ is

$$\sum_{\Pi}\int_{\mathcal{S}^2}|g_{i,\hat{k}\Pi}|^2\mathrm{d}^2\hat{\boldsymbol{k}} = d_{ie}{}^2\int_{\mathcal{S}^2}\sum_{\Pi}|\hat{\boldsymbol{d}}_{ie}\cdot\hat{\boldsymbol{\epsilon}}_{\hat{k}\Pi}|^2\,\mathrm{d}^2\hat{\boldsymbol{k}} = \frac{8\pi}{3}d_{ie}{}^2. \tag{2.72}$$

This leaves the terms in k and the integration over $\mathrm{d}t'$, for which we make a Markov-type approximation [Mil76] by replacing $\tilde{c}_e(t') \approx \tilde{c}_e(t)$ in the integrand, such that for times t long compared to the optical timescale $\omega_{ei}{}^{-1}$

$$\frac{\mathrm{d}}{\mathrm{d}t}\tilde{c}_e(t) \approx -\sum_{i\in I_L}\frac{8\pi}{3}d_{ie}{}^2\int_0^\infty k^2\tilde{g}_k^2\underbrace{\int_0^t e^{i(\omega_{ei}-kc)(t-t')}\,\mathrm{d}t'}_{\approx\pi\delta(\omega_{ei}-kc)-i\mathcal{P}(\frac{1}{\omega_{ei}-kc})}\,\mathrm{d}k\,\tilde{c}_e(t)$$

$$\approx^{22} -\sum_{i\in I_L}\frac{2(2\pi)^2}{3c}d_{ie}{}^2\,k_{ei}{}^2\,\tilde{g}_{k_{ei}}^2\,\tilde{c}_e(t), \tag{2.73}$$

where we have written $k_{ei} := \omega_{ei}/c$ for the angular wavenumber exactly corresponding to the energy difference between the atomic levels. Hence, the excited state decays exponentially:

$$\tilde{c}_e(t) = e^{-\frac{\Gamma}{2}t}, \quad \Gamma := \sum_{i\in I_L}A_{ei}, \quad A_{ei} := \frac{k_{ei}{}^3 d_{ie}{}^2}{3\pi\hbar\varepsilon_0}. \tag{2.74}$$

---

[22] $\mathcal{P}$ in the second term of the approximation to the $\mathrm{d}t'$ integral denotes the Cauchy principal value. As the term is purely imaginary, it leads to a divergent level shift. It can be recognised as part of the Lamb shift after renormalisation [Mil76], but we simply drop it here, considering any level shifts to be already included in the atomic energy levels $\hbar\omega_i$.

This is the well-known expression for the decay constant Γ of the excited state population $|\tilde{c}_e(t)|^2$ as the sum of all Einstein coefficients for the individual transitions A_{ei}, where we can trace back the k_{ei}^3 dependence to the density of e.m. modes in vacuum.

Photon wavepackets entangled with lower ion states. Inserting eq. (2.74) back into the mode-continuum equivalent of eq. (2.68) and integrating, we obtain

$$\tilde{c}_{i,k\Pi}(t) = \tilde{g}_k\, g_{i,\hat{k}\Pi}\, \frac{1 - e^{-\left(\frac{\Gamma}{2} + i(\omega_{ei} - kc)\right)t}}{\frac{\Gamma}{2} + i(\omega_{ei} - kc)} \tag{2.75}$$

for the individual mode components, and we can write (leaving the interaction picture) the joint ion–photon state as

$$
\begin{aligned}
|\psi(t)\rangle &= e^{-\frac{\Gamma}{2}t} e^{-i\omega_e t}\, |e\rangle \otimes |0\rangle \\
&+ \sum_{i \in I_L} e^{-i\omega_i t}\, |i\rangle \otimes \left(\sum_{\Pi} \int_{\mathbb{R}^3} \tilde{g}_k\, g_{i,\hat{k}\Pi}\, \frac{1 - e^{-\left(\frac{\Gamma}{2} + i(\omega_{ei} - kc)\right)t}}{\frac{\Gamma}{2} + i(\omega_{ei} - kc)} e^{-ikct} a_{k\Pi}^\dagger\, d^3 k \right) |0\rangle .
\end{aligned}
\tag{2.76}
$$

The bracketed expression defines a wavepacket corresponding to each decay channel in terms of its plane-wave expansion. In the photon wavefunction formalism, we can represent the wavepacket corresponding to an $e \mapsto i$ decay by an explicit mode function (cf. eq. (2.35))

$$\boldsymbol{f}_i(\boldsymbol{r}, t) = i\sqrt{\frac{\hbar c}{2\varepsilon_0}} N_i^{-1} \sum_{\Pi} \int_{\mathbb{R}^3} \sqrt{k}\, \tilde{g}_k\, g_{i,\hat{k}\Pi}\, \frac{1 - e^{-\left(\frac{\Gamma}{2} + i(\omega_{ei} - kc)\right)t}}{\frac{\Gamma}{2} + i(\omega_{ei} - kc)} \hat{\boldsymbol{\epsilon}}_{\hat{k}\Pi} e^{ik(\hat{k}\cdot\boldsymbol{r} - ct)} d^3 k, \tag{2.77}$$

where $0 < N_i < 1$ is a yet to be determined, dimensionless normalisation constant such that $\boldsymbol{f}_i(\boldsymbol{r}, t)$ contains exactly one photon[23]. This way, we can for each decay channel define new creation and annihilation operators $\{\tilde{a}_i^\dagger, \tilde{a}_i\}$ such that the state and the electric field operator become

$$
\begin{aligned}
|\psi(t)\rangle &= e^{-\left(\frac{\Gamma}{2} + i\omega_e\right)t} |e\rangle \otimes |0\rangle + \sum_{i \in I_L} N_i\, e^{-i\omega_i t} |i\rangle \otimes \tilde{a}_i^\dagger |0\rangle , \\
\boldsymbol{E}_\perp(\boldsymbol{r}, t) &= \sum_{i \in I_L} \boldsymbol{f}_i(\boldsymbol{r}, t)\, \tilde{a}_i + \text{h. c..}
\end{aligned}
\tag{2.78}
$$

Note that the modes associated with different decay channels will not in general be orthogonal.

---

[23] Some care is needed as the coefficients explicitly depend on t in the numerator, not only on k. This contribution decays exponentially for $t \gg \Gamma^{-1}$, though, and only serves a formal role in suppressing an incoming-wave contribution in eq. (2.81).

To explicitly compute the mode functions, we again split up the integral into radial and angular parts,

$$
\begin{aligned}
\boldsymbol{f}_i(\boldsymbol{r},t) &= \mathrm{i}\sqrt{\frac{\hbar c}{2\varepsilon_0}}\, N_i^{-1} \sum_\Pi \int_{\mathbb{R}^3} k \sqrt{\frac{c}{2(2\pi)^3 \hbar \varepsilon_0}}\, d_{ie}\hat{\boldsymbol{d}}_{ie}\cdot\hat{\boldsymbol{\epsilon}}_{\hat{\boldsymbol{k}}\Pi}\,\frac{1-\mathrm{e}^{-\left(\frac{\Gamma}{2}+\mathrm{i}(\omega_{ei}-kc)\right)t}}{\frac{\Gamma}{2}+\mathrm{i}(\omega_{ei}-kc)}\,\hat{\boldsymbol{\epsilon}}_{\hat{\boldsymbol{k}}\Pi}\,\mathrm{e}^{\mathrm{i}k(\hat{\boldsymbol{k}}\cdot\boldsymbol{r}-ct)}\mathrm{d}^3\boldsymbol{k} \\
&= \mathrm{i}\frac{d_{ie}\,c}{2\varepsilon_0(2\pi)^{3/2}}\int_0^\infty k^3\,\frac{1-\mathrm{e}^{-\left(\frac{\Gamma}{2}+\mathrm{i}(\omega_{ei}-kc)\right)t}}{\frac{\Gamma}{2}+\mathrm{i}(\omega_{ei}-kc)}\,\mathrm{e}^{-\mathrm{i}kct}\int_{\mathcal{S}^2}\mathrm{e}^{\mathrm{i}k(\hat{\boldsymbol{k}}\cdot\boldsymbol{r})}\underbrace{\sum_\Pi\left(\hat{\boldsymbol{d}}_{ie}\cdot\hat{\boldsymbol{\epsilon}}_{\hat{\boldsymbol{k}}\Pi}\right)\hat{\boldsymbol{\epsilon}}_{\hat{\boldsymbol{k}}\Pi}}_{=\hat{\boldsymbol{d}}_{ie}-(\hat{\boldsymbol{d}}_{ie}\cdot\hat{\boldsymbol{k}})\hat{\boldsymbol{k}}}\mathrm{d}^2\hat{\boldsymbol{k}}\,\mathrm{d}k.
\end{aligned}
\tag{2.79}
$$

The angular integral can be straightforwardly evaluated by choosing spherical coordinates with the polar axis aligned with $\hat{\boldsymbol{r}}$ (the normalised position coordinate, $\boldsymbol{r}=r\,\hat{\boldsymbol{r}}$):

$$
\begin{aligned}
\int_{\mathcal{S}^2}\mathrm{e}^{\mathrm{i}k(\hat{\boldsymbol{k}}\cdot\boldsymbol{r})}\left(\hat{\boldsymbol{d}}_{ie}-(\hat{\boldsymbol{d}}_{ie}\cdot\hat{\boldsymbol{k}})\hat{\boldsymbol{k}}\right)\mathrm{d}^2\hat{\boldsymbol{k}} &= 2\pi\,\hat{\boldsymbol{d}}_{ie}\cdot\underbrace{\int_0^\pi \mathrm{e}^{\mathrm{i}kr\cos\theta}\left(\mathbb{1}-(\sin\theta)^2\frac{\mathbb{1}-\hat{\boldsymbol{r}}\hat{\boldsymbol{r}}}{2}+(\cos\theta)^2\hat{\boldsymbol{r}}\hat{\boldsymbol{r}}\right)\sin\theta\,\mathrm{d}\theta}_{=\frac{\mathrm{e}^{\mathrm{i}kr}}{kr}\left((-\mathrm{i}+\frac{1}{kr}+\frac{\mathrm{i}}{(kr)^2})(\mathbb{1}-\hat{\boldsymbol{r}}\hat{\boldsymbol{r}})+(-\frac{1}{kr}-\frac{\mathrm{i}}{(kr)^2})\hat{\boldsymbol{r}}\hat{\boldsymbol{r}}\right)+\mathrm{c.\,c.}} \\
&= -2\pi\mathrm{i}\left(\frac{\mathrm{e}^{\mathrm{i}kr}}{kr}-\frac{\mathrm{e}^{-\mathrm{i}kr}}{kr}\right)\left(\hat{\boldsymbol{d}}_{ie}-(\hat{\boldsymbol{d}}_{ie}\cdot\hat{\boldsymbol{r}})\hat{\boldsymbol{r}}\right)+\mathrm{O}\left(\frac{1}{(kr)^2}\right)
\end{aligned}
\tag{2.80}
$$

This exactly mirrors the form of classical dipole radiation; neglecting the terms in higher powers of $(kr)^{-1}$ corresponds to considering the far-field limit $kr\gg 1$. To evaluate the remaining integral over k, we again assume that the emission is narrow-band ($\Gamma\ll\omega_{ei}$) such that we can approximate the k^2 factor by $(\omega_{ei}/c)^2$ and in excellent approximation extend the lower integration bound to $-\infty$. Analytically continuing the integrand into the complex plane and integrating along a semicircular contour then gives

$$
\int_0^\infty k^2\,\frac{1-\mathrm{e}^{-\left(\frac{\Gamma}{2}+\mathrm{i}(\omega_{ei}-kc)\right)t}}{\frac{\Gamma}{2}+\mathrm{i}(\omega_{ei}-kc)}\,(\mathrm{e}^{\mathrm{i}kr}-\mathrm{e}^{-\mathrm{i}kr})\mathrm{e}^{-\mathrm{i}kct}\mathrm{d}k \approx \frac{2\pi\omega_{ei}^2}{c^3}\mathrm{e}^{-\left(\frac{\Gamma}{2}+\mathrm{i}\omega_{ei}\right)\left(t-\frac{r}{c}\right)}\Theta\left(t-\frac{r}{c}\right),
\tag{2.81}
$$

where Θ is the Heaviside unit step function (with $\Theta(x)=1$ for $x>0$ and 0 otherwise).

We have thus obtained an explicit expression for the mode functions from eq. (2.78); a spontaneous decay from $|e\rangle$ to $|i\rangle$ is associated with a photon wavepacket

$$
\boldsymbol{f}_i(\boldsymbol{r},t) = N_i^{-1}\frac{d_{ie}\,\omega_{ei}^2}{2\sqrt{2\pi}\,\varepsilon_0\,c^2}\frac{\mathrm{e}^{-\mathrm{i}\omega_{ei}(t-\frac{r}{c})}}{r}\left(\hat{\boldsymbol{d}}_{ie}-(\hat{\boldsymbol{d}}_{ie}\cdot\hat{\boldsymbol{r}})\hat{\boldsymbol{r}}\right)\mathrm{e}^{-\frac{\Gamma}{2}\left(t-\frac{r}{c}\right)}\Theta\left(t-\frac{r}{c}\right),
\tag{2.82}
$$

describing a radial wave front $\mathrm{e}^{-\mathrm{i}\omega_{ei}t'}/r$ emanating from the atom in the origin of the coordinate system with a temporal envelope $\mathrm{e}^{-\Gamma t'/2}$ for retarded times $t':=t-r/c>0$. Its angular amplitude and polarisation pattern is given by the classical dipole radiation pattern $\hat{\boldsymbol{d}}_{ie}-(\hat{\boldsymbol{d}}_{ie}\cdot\hat{\boldsymbol{r}})\hat{\boldsymbol{r}}=-\hat{\boldsymbol{r}}\times(\hat{\boldsymbol{r}}\times\hat{\boldsymbol{d}}_{ie})$.

Finally, we can choose the normalisation factor N_i such that each mode contains, for long times $t \rightarrow \infty$, exactly one photon[24], i.e. the overall electrical energy $\int_{\mathbb{R}^3} \frac{\varepsilon_0}{2} |f_i(r, t)|^2 \mathrm{d}^3 r$ is exactly $\frac{\hbar \omega_{ei}}{2}$. This gives $N_i{}^2 = A_{ei}/\Gamma$, and

$$|\psi(t)\rangle \overset{t \gg \Gamma^{-1}}{\approx} \sum_{i \in I_L} \sqrt{\frac{A_{ei}}{\Gamma}} \, \mathrm{e}^{-\mathrm{i}\omega_i t} \, |i\rangle \otimes \tilde{a}_i^\dagger |0\rangle \,. \qquad (2.83)$$

The probability to find the ion in each of the lower states at long times is given by the fractional contribution of the respective Einstein A coefficient to the total linewidth Γ.

2.5   Heralded remote entanglement generation

Per eq. (2.55), the interaction between the atomic dipole moment[25] and the electric field is unitary, even in free space. Consequently, the simplest way to photonically mediate an entangling interaction between two ions might appear to be for one ion to spontaneously emit a photon entangled with its internal state, and for the other to absorb it, leaving the ions in an entangled state. Such a direct transmission scheme is in theory indeed conceivable, but would be exceedingly challenging to achieve in practice, as any non-idealities would affect the state transfer (and hence the quality of the generated entanglement).

Two issues are immediately apparent: First, as dipole emission covers the full 4π of solid angle, reaching the "strong coupling" regime in free space would require extremely high-numerical-aperture optics that can collect all this emission and re-focus it at the location of the receiver ion. Secondly, as the absorption process would be the time inverse of the emission process, for it to occur with unit probability, the temporal shape of the photon would also need to be flipped. Of course, this is trivial to ensure if the photon time envelope is symmetrical, but far from that for the exponentially decaying amplitude associated with spontaneous decay (and consequently, the exponentially rising amplitude necessary for efficient free-space absorption). This has indeed been verified in experiments with attenuated laser pulses [Alj+13; Leo+16], but losslessly inverting the temporal shape of narrow-band photons is challenging, requiring e.g. the use of specifically tailored optical cavities [Sri+14].

---

[24]This can be viewed as an explicit step to assume the perspective outlined at the beginning of this chapter, where we separate the photon generation process from the further operations. In this setting, it is convenient to include the exponential envelope in $f_i(r, t)$ rather than explicitly in the time-dependent state vector. Coincidentally, in our experiment the optical path to the single-photon detectors is not long enough for emission and detection processes actually to be well-separated in time, but as we do not interact with the ions before $t \gg \Gamma^{-1}$, we cannot observe anything outside the ground-state subspace. If, hypothetically, we were able to rapidly measure the ion state, the exponentially decaying amplitude from eq. (2.74) would not be a good approximation; a photon detection projects the ion into one of the ground states with certainty.

[25]Other types of transitions could in theory also be employed for remote entanglement generation, though electric dipole (E1) transitions are the most obvious candidates, as their coupling strength tends to be the largest.
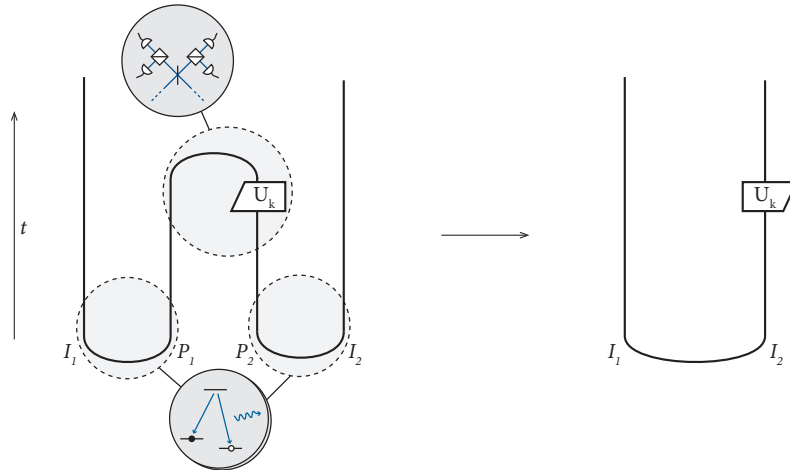
Figure 2.4: Heralded entanglement generation between remote qubits across a photonic link. Creating two ion–photon maximally entangled states and performing a Bell-basis measurement on the photons P_1 and P_2 leaves the ions I_1 and I_2 in a maximally entangled state. This is visually evident in the graphical formalism of Coecke et al. [Coe10] employed here, where maximally entangled states are represented by cups ($\bigcup$), Bell-basis measurements by caps ($\bigcap$, with a unitary correction U_k depending on the outcome k of the measurement), and "yanking" lines straight results in valid mathematical equations.

Placing the atomic emitters and absorbers in optical cavities can in principle help address both of these challenges. The modified mode structure of the electromagnetic vacuum between the cavity mirrors can greatly enhance the coupling to the well-defined cavity mode compared to isotropic free space, and by using the cavity-enhanced emission as one leg of a Λ-system Raman process, where the other leg is driven by an laser pulse, the coupling can be modulated in time to ensure efficient reabsorption. In this framework of cavity quantum electrodynamics, quantum teleportation and remote qubit entanglement via direct photon transmission have indeed been demonstrated [Kim08; RR15], but doing so at a high level of performance (transmission/entanglement fidelity) is still far from trivial. For trapped ions in particular, the interaction between the dielectric cavity mirrors and the ion charge makes them technically harder to implement yet (cf. §§6.4 and 8.1.3).

Even if all this could theoretically be achieved, practical realisations will always suffer from imperfections. In particular, for long-distance quantum networking application, a considerable amount of link loss is unavoidable. Fortunately, strong coupling and deterministic interaction turn out not to actually be necessary, though. The key observation is that a process that only succeeds a fraction of the time[26] is perfectly serviceable as long as there is some signal to indicate whether it in fact did, such that repeated entanglement attempts can be made until a success is eventually observed. Such procedures are known as "heralded" entanglement generation, or "event-ready", schemes.

One such heralded entanglement generation scheme, which is in ubiquitous use across virtually all qubit platforms, is depicted in fig. 2.4. By bringing together one photon from

---

[26]Or, in the unitary picture: [...] a process where the state component corresponding to successful generation of entanglement only has a small amplitude [...]

each node on a central beamsplitter and placing detectors in its output, a probabilistic Bell-basis measurement can be applied to the photons. Such a projective measurement, the success of which is heralded by two simultaneous clicks in the photon detectors, applied to a pair of ion–photon Bell states, itself leaves the ions in a maximally entangled, pure state – a process known as "entanglement swapping" [Żuk+93]. It is remarkable that no information travels directly from one node to the other in this entanglement generation process (only the heralding success signal from the central station to each node).

The realisation of a Bell state measurement in linear optics relies solely on the indistinguishability of the photons; in the simplest case of a 50 : 50 beamsplitter directly followed by two detectors, the Hong–Ou–Mandel effect [HOM87] ensures that coincident clicks project the photons into the one antisymmetric Bell state $|\Psi^-\rangle$. Thus, the chosen photon basis states do not necessarily need to be amenable to direct manipulation as a qubit, as long as maximally entangled ion–photon states are quick to prepare with high fidelity, easy to collect with good efficiency, and identical across time and different nodes. In the experiments described here, however, we will employ photons entangled with the ion in their polarisation, which is readily manipulated. Not only does this allow us to analyse the behaviour of individual nodes in isolation, but with the addition of two polarising beamsplitters and two (for a total of four) single-photon detectors, two out of four Bell states (doubling the heralding probability over a simple two-detector measurement). As the successful projection onto a Bell state requires the detection of two photons, the success probability is quadratic in the individual collection/detection probabilities for a photon from each node (up to an asymptotic limit of $\frac{1}{2}$ in the case of linear optics, where only two out of four Bell states can be distinguished). A practically very useful property of such schemes is also that the path length stability requirements for the optical link are not given by the optical wavelength, but only the photon overlap condition and any frequency splitting between the qubit states, and thus significantly more lenient (as discussed further in chapter 4).

Finally, it should be noted that heralded remote entanglement of this nature can also be realised in schemes similarly exploiting the interference on a beamsplitter, but heralded through the detection of only a single photon (where the entanglement mechanism can be though of as stemming from the ambiguity about which node the detected photon originated from) [Cab+99]. Such schemes, which can equivalently be considered to use the number-state basis as the photon qubit encoding, require the optical path linking the nodes to be interferometrically stable, and (for the production of single Bell pairs [CB08]), have an intrinsic rate–fidelity trade-off. While they have the advantage of only depending linearly on the one-node photon collection probability in their success rate, the sensitivity to wavelength-scale path fluctuations is an additional challenge in trapped ions. Such schemes have attracted attention in some other qubit technologies, but have received little attention in trapped ions (with the exception of ref. [Slo+13]); we will not consider them

here any further.

In chapter 4, we will analyse a two-photon heralded remote entanglement scheme in great detail. First, however, we turn our attention to the concrete experimental platform used for our remote entanglement demonstrations.

3 Alice and Bob: Twin ion-trap apparatus

A key component of the state-of-the-art, stable remote entanglement results to be discussed in chapters 6 and 7 was the construction of two trapped-ion quantum network nodes. Compared to the traditional way of constructing experiments in atomic and molecular physics (centred around flexible optical breadboards, etc.), many of the subsystems were designed and engineered intentionally with the requirements of few-qubit remote entanglement experiments in mind, with a particular focus on robustness and maintainability. While the design and construction of this system has taken up significant effort over the span of this project, due to space constraints I will only give a very brief summary of the apparatus here to provide some context for the rest of this thesis. The interested reader is referred to L. Stephenson's recent DPhil thesis for a more complete description [Ste19].

3.1 The ions: $^{43}Ca^+$ and $^{88}Sr^+$

In typical trapped-ion quantum information experiments, one or more positively charged atomic ions are confined in radio-frequency Paul traps under ultra-high vacuum vacuum
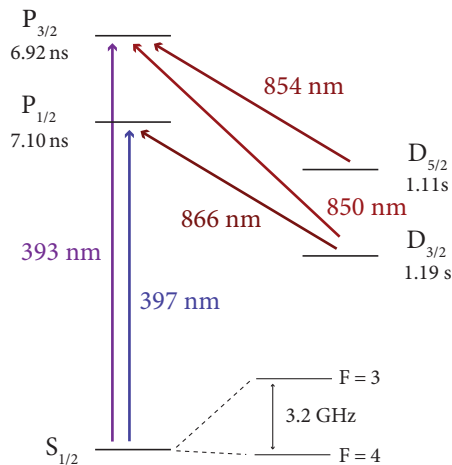
Figure 3.1: Diagram of relevant energy levels in $^{43}Ca^+$ (not to scale), with transitions addressed by lasers in the experiment. Hyperfine splitting for P and D levels and Zeeman splittings are not shown. (Data from refs. [JC93; Sah+06].)
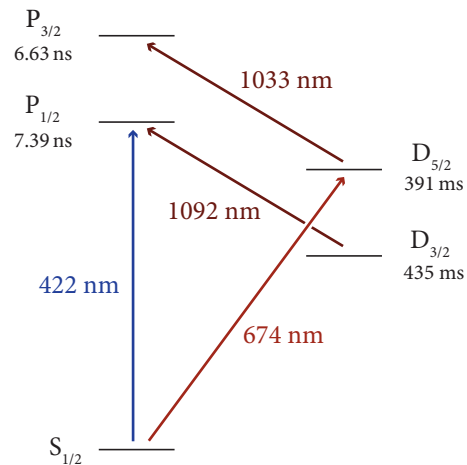
Figure 3.2: Diagram of relevant energy levels in $^{88}Sr^+$ (not to scale), with transitions addressed by lasers in the experiment. Zeeman splittings are not shown. (Data from refs. [San12; Let+05; Man+99; PBL95].)

(ideally well below 10^{-9} Pa to avoid ion loss from background gas collisions). A comprehensive toolbox of techniques using lasers and radio-frequency radiation to manipulate the system exists, including isotope-selective photoionisation, multiple forms of laser cooling, dissipative preparation of the internal electronic state, and coherent single- and multi-qubit operations. These are well explored theoretically and have been implemented in a number of atomic species [Win+98; Lei+03].

A main focus of our group in recent years was the development of techniques for the manipulation of ^{43}Ca$^+$ (see fig. 3.1). Its nuclear spin of $\frac{7}{2}$ somewhat complicates laser cooling, but qubits stored in the hyperfine ground state manifold (splitting $\approx$ 3.2 GHz) are among the best-performing in any technology. High-fidelity state preparation and measurement (error $\epsilon = 7 \times 10^{-4}$), single-qubit gates ($\epsilon = 1 \times 10^{-6}$), and long coherence times ($T_2^* > 50$ s) have been demonstrated [Har+14], as well as high-fidelity laser-driven single-qubit ($\epsilon = 7 \times 10^{-5}$) and two-qubit gates ($\epsilon = 1 \times 10^{-3}$) [Bal+16]. At the same time, the required laser wavelengths are all near the visible regime, where commercial diode lasers and off-the-shelf optical components and fibres are readily available.

Including a second ion species with different transition wavelengths provides an intrinsic means to target only a subset of ions in a trap, for example to cool the motion of an ion string while keeping information stored on another species (sympathetic cooling). For these experiments, we have chosen ^{88}Sr$^+$ as a companion species to ^{43}Ca$^+$ (see fig. 3.2). It has no hyperfine structure, which simplifies cooling and is beneficial for achieving a high rate of ion–photon entanglement. The transition wavelengths are well-separated from ^{43}Ca$^+$, yet close enough to make common optical systems for both species straightforward to design. The mass ratio of $m_{\mathrm{Sr}}/m_{\mathrm{Ca}} \approx 2$ is also moderate enough not to pose any challenges for trap stability, although the coupling of the radial modes between the species is already weak enough to only realistically[1] leave the axial modes as an attractive choice for multi-qubit gates [Hom13]. We have recently demonstrated such a ^{43}Ca$^+$–^{88}Sr$^+$ mixed-species gate in a different apparatus (not described here, see ref. [Thi19]), achieving a Bell-state fidelity of $\mathcal{F} = 99.8(1)$ % [Hug+20]. The experiments described in this thesis involve only a single ^{88}Sr$^+$ ion; we will only mention the accommodations for ^{43}Ca$^+$ in passing here[2].

---

[1]The secular oscillation frequency of a single ion in the rotating pseudopotential of a Paul trap is inversely proportional to the ion mass. As a consequence, calculations for e.g. the simplest ^{43}Ca$^+$–^{88}Sr$^+$ crystal in a trap with single-^{43}Ca$^+$ frequencies {2.0 MHz, 8.5 MHz, 9.5 MHz} give the amplitude contribution of one species to the eigenvector of the radial mode dominated by the other to only 2 % (in terms of mass-weighted positions, i. e. as it enters the effective Lamb–Dicke parameter describing the strength of laser-motion coupling).

[2]The complete system has been constructed, however, and at the time of writing, ^{43}Ca$^+$–^{43}Ca$^+$, ^{88}Sr$^+$–^{43}Ca$^+$ and ^{88}Sr$^+$–^{88}Sr$^+$ two-qubit gates have been achieved.
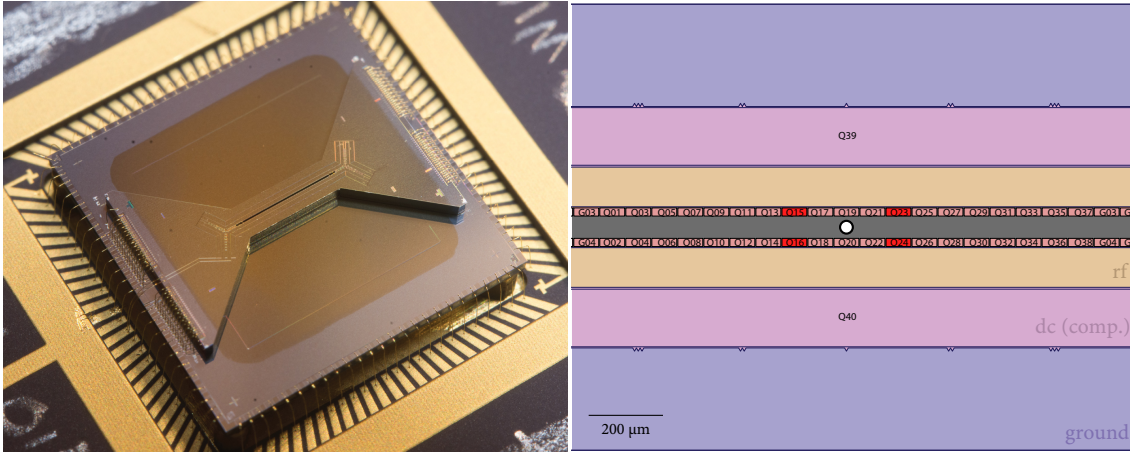
Figure 3.3: The Sandia *High Optical Access 2* (HOA2) microfabricated surface-electrode trap. *Left:* The trap consists of a linear central region with two Y-shaped junctions on either end. A slot through the substrate provides optical access from the back side; the bow-tie shape provides a large free solid angle for beams skimming the trap. *Right:* Detail of the electrode arrangement in the middle of the central isthmus, close to the design ion position marked by a white dot (design null 70 μm above the surface). The oblong electrodes Q39 and Q40 are used for stray field compensation, the four highlighted DC electrode segments Q15, Q16, Q23, Q24 for link disconnection in chapter 7. (Photograph and original drawing courtesy of Sandia National Laboratories.)

3.2 Apparatus overview

The experiment consists of two trapped-ion nodes that are (mirror-symmetric) copies of each other, and almost completely independent except for a shared set of laser sources. A high-level block diagram of the involved subsystems is shown in fig. 3.4, some of which are briefly described in the following; the reader is once again referred to the thesis of L. Stephenson [Ste19] for details.

3.2.1 Ion traps

The ion trap nodes are built around the "High Optical Access 2.1" (HOA2) microfabricated surface-electrode trap manufactured by Sandia National Laboratories [Mau16] (see fig. 3.3). It was chosen as a proven design, and its 94 independent control voltages provide more than sufficient flexibility to establish multiple independent trap zones for the demonstration of more complex algorithms (see outlook in §8.2). The design ion height is 68 μm above the surface, with segmented DC control electrodes allowing the independent control of several ion strings (see fig. 3.3). The RF electrodes are resonantly matched[3] to 50 Ω at 51 MHz with a quality factor $Q \sim 50$, resulting in a zero-peak RF amplitude of ~ 170 V. The

---

[3]In Alice, a single-stage lumped-element *LC* matching network is used, in Bob, a toroidal transformer. The exact step-up factor depends on parasitic losses and the tuning on various parasitic reactances; a higher voltage somewhat closer to the rated trap voltage limit of 300 V, at a slightly increased RF drive frequency, would be desirable to reach lower motional heating rates.

Figure 3.4: High-level component diagram of the remote networking apparatus. With the exception of the laser sources (and their diagnostic systems), each of the nodes Alice and Bob has its independent copy of the necessary systems. Figure adapted from ref. [Ste19].
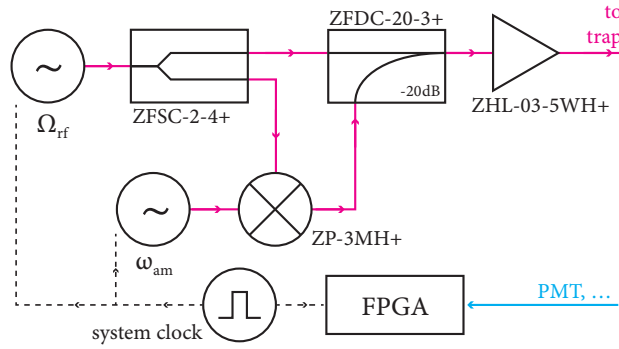
Figure 3.5: Block diagram of the RF generation circuit (part numbers for suitable components by Mini-Circuits, Inc.). A DDS channel at Ω_{rf} supplies the unmodulated trapping potential (optionally through a limiting amplifier for further passive amplitude stabilisation). To provide a small, tunable amount of amplitude modulation without affecting the mode stability in regular operation, the signal is split into two paths using a passive power splitter, one of which is modulated using a double-balanced diode-ring mixer, whose IF port is driven from the second Urukul channel set to ω_{AM}. The signals are then recombined using a directional coupler and amplified before the trap resonator. In this configuration, the phase of both the Ω_{RF} and ω_{AM} tones is stable and known in absolute terms with respect to the FPGA timeline, enabling simple demodulation of the PMT photon arrival times in the digital domain for micromotion compensation.

DC electrodes are driven by a multi-channel 16-bit, $\pm 10\,\mathrm{V}$ digital-to-analogue converter[4] (DAC), with voltages chosen to result in typical mode frequencies of 1.85 MHz, 2.28 MHz, and 2.90 MHz for the axial and two radial modes (when combined with the RF pseudopotential). A static quadrupole term is added such that the radial mode vectors are rotated by 28° about the trap axis, resulting in overall angles of 70° and 51° between the cooling beams and the low- and high-frequency radial modes, respectively.

For the stray field compensation method described in chapter 5, RF amplitude modulation is applied using a mixer circuit (fig. 3.5) driven by a direct digital synthesis source (DDS) with phase referenced to the same clock used for input timestamping (fluorescence photon timestamping). The sensitivity of the electric field at the ion position to the voltages applied to the global compensation electrodes in the in-plane and out-of-plane directions is $0.5\,\mathrm{Vm^{-1}/mV}$ and $0.6\,\mathrm{Vm^{-1}/mV}$, respectively, such that the DAC resolution limits the stray field compensation precision to $0.2\,\mathrm{Vm^{-1}}$.

3.2.2   Non-coherent/dissipative operations

All ion state manipulation is performed through laser beams which propagate parallel to the trap plane and are focussed at the trap centre (fig. 3.6).

For $^{88}\mathrm{Sr}^{+}$, atoms are loaded from a thermal source [Bal+18] through a two-stage resonant photoionisation process [Luc+04] using laser beams at 461 nm and 378 nm. The $5s\,S_{1/2} \leftrightarrow 5p\,P_{1/2} \leftrightarrow 4d\,D_{3/2}$ Doppler cooling cycle is addressed by 422 nm and 1092 nm

---

[4]The *Zotino* module from the ARTIQ/Sinara open-source hardware ecosystem; later replaced by *Fastino* modules for faster update rates.
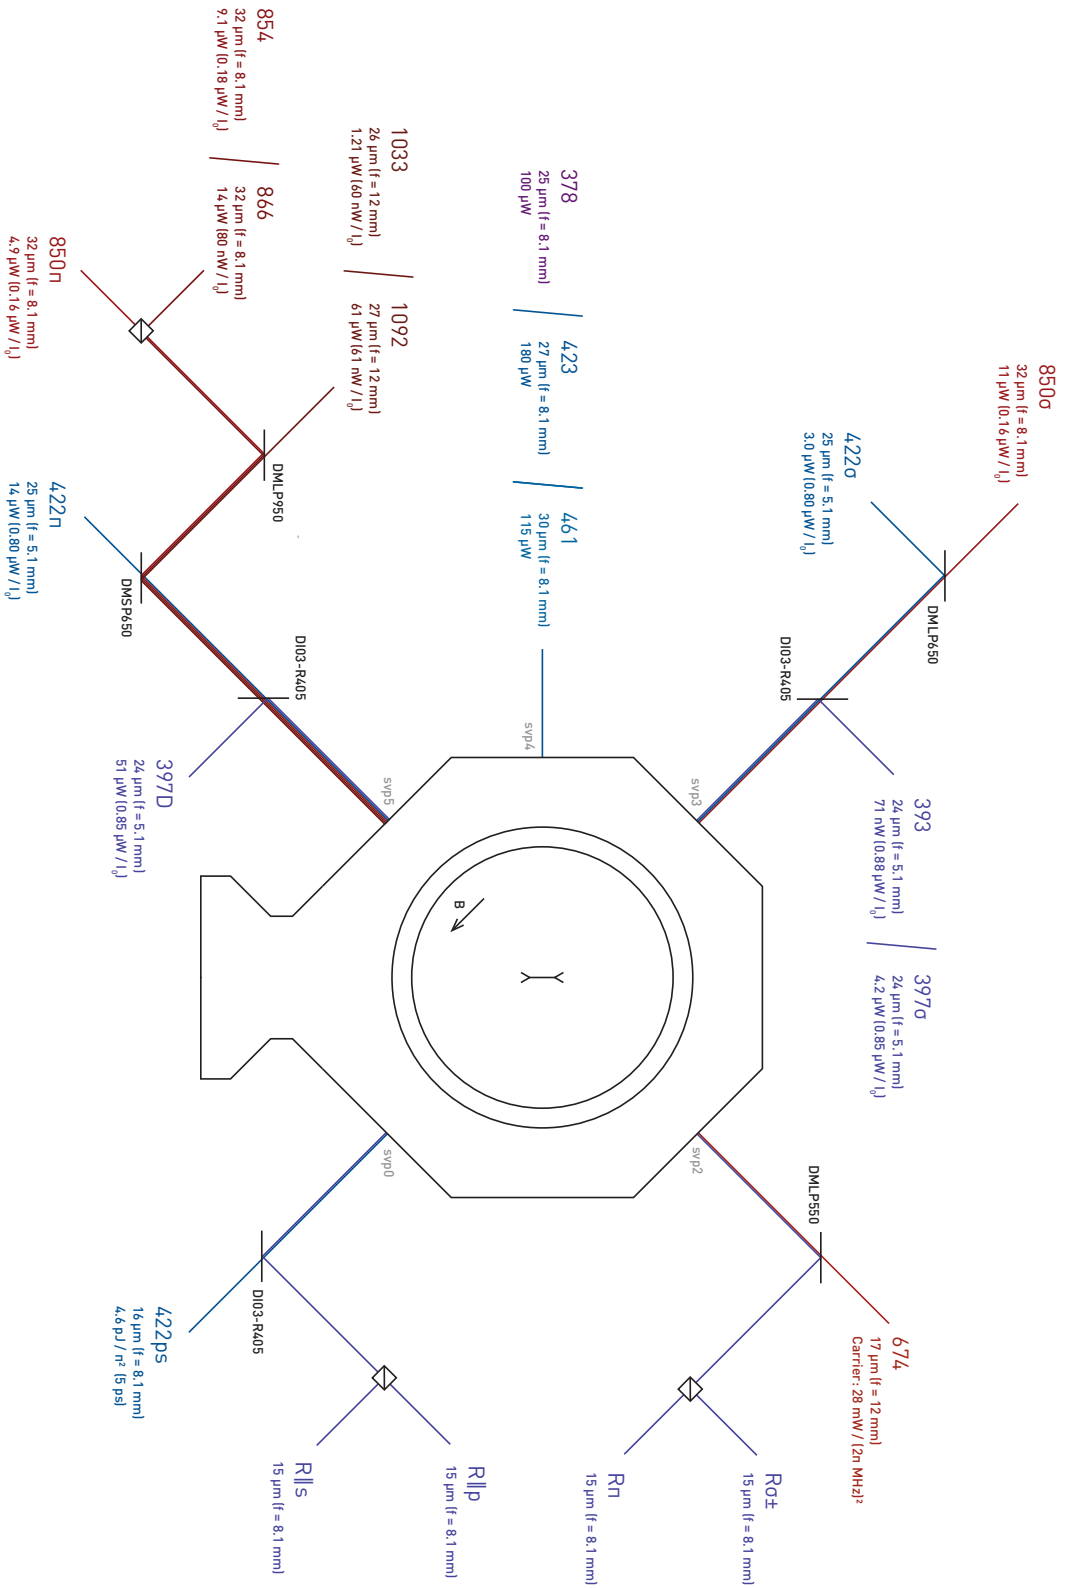
Figure 3.6: Schematic layout of all laser beams entering the trap vacuum systems (combined by polarising beamsplitters and dichroic mirrors), labelled by their respective wavelengths in nm (σ denotes purely circularly polarised beams, R the beams near 400 nm for Raman interactions in $^{43}\text{Ca}^+$). Spot sizes are estimates for the $\frac{1}{e^2}$-radius assuming aberration-free Gaussian beams and nominal focal lengths. I_0 denotes the saturation intensity, where the transition rate coincides with the Clebsch–Gordan-weighted A-coefficient. Budgeted powers do not include losses. The arrangement is as shown for Bob when viewing the trap from the top side along the high-NA imaging system; that of Alice is its mirror image.

beams (of approximately π and $\sigma^{\pm}$ polarisations) perpendicular to the 0.5 mT magnetic field, each oriented at 45° to the trap axis. To prepare the ion in $\left|5s\,{}^2S_{1/2}, m_J = -\frac{1}{2}\right\rangle$ though optical pumping, a σ^--polarised 422 nm beam is applied (together with the 1092 nm beam to avoid population trapped in the metastable states). State detection also uses the 422 nm and 1092 nm cooling beams; if the ion has been coherently transferred to the $D_{5/2}$ level using 674 nm laser pulses (see §3.2.3), it remains dark, while it fluoresces bright otherwise. A 1033 nm laser dissipatively returns $D_{5/2}$ population to the cooling cycle.

Ion fluorescence is collected through a 60 μm slot in the trap centre (numerical aperture 0.25) and imaged onto a photomultiplier tube (PMT), resulting in an overall collection efficiency of $\eta \approx 0.4\%$ with very low background from laser scatter. Typical signal count rates for a single bright ion are 160 kHz vs. < 100 Hz background, leading to negligible statistical state discrimination error for typical readout times of 200 μs.

^{43}Ca$^+$ is similarly loaded from a thermal source, ionised using 423 nm and 378 nm lasers, and cooled and initialised by addressing the $4s\,{}^2S_{1/2} \leftrightarrow 4p\,{}^2P_{1/2}$ and $4p\,{}^2P_{1/2} \leftrightarrow 3d\,{}^2D_{3/2}$ transitions using 397 nm and 866 nm lasers. In ^{43}Ca$^+$, we opt to avoid the requirement for a narrow-linewidth laser by state-selectively shelving the ion to the $D_{5/2}$ level via the 393 nm and 850 nm excitation [Mye+08]; the 854 nm laser returns shelved population to the cooling cycle.

Most laser sources are external cavity diode lasers[5], stabilised to piezo-actuated reference cavities using the Pound–Drever–Hall locking scheme [Dre+83] for 393 nm, 397 nm, 422 nm, 866 nm and 1092 nm, and digitally stabilised with a low feedback bandwidth to a Fizeau-interferometer-based wavemeter[6] for 423 nm, 461 nm, 854 nm and 1033 nm. The 378 nm laser is a free-running diode laser[7]. The beams are frequency-shifted and switched as required using acousto-optical modulators (AOMs). After delivery to the trap nodes through optical fibres, the beams are filtered in polarisation and stabilised in power by sampling a small portion of the beam onto a photodiode and feeding back onto the RF amplitude driving a respective AOM. The exception to this are the photoionisation lasers, which are controlled via mechanical shutters.

3.2.3 *Coherent qubit operations*

The currents through three copper wire coils surrounding the trap vacuum system are set such that the static magnetic field at the position of the ion is $B = 0.5000$ mT. The precise value is an arbitrary choice, but a field of roughly this magnitude is large enough to give appreciable (~ 10 MHz) Zeeman splittings, but still small enough to retain a low

---

[5] Toptica *DL Pro* with *DLC Pro* controllers.
[6] HighFinesse *WS-7*.
[7] Toptica *iBeam smart*.

magnetic field sensitivity of the $|F = 4, m_F = 0\rangle \leftrightarrow |F = 3, m_F = 0\rangle$ zero-field clock qubit in the $4s\,^2S_{1/2}$ ^{43}Ca$^+$ ground state[8].

At this magnetic field flux density, the splitting between the $|S_{1/2}, m_J = \pm 1/2\rangle$ ground state Zeeman qubit levels in ^{88}Sr$^+$, where the ion–photon entanglement is naturally created, is 14.01 MHz. Instead of directly manipulating the qubit state with an RF magnetic field, we choose to manipulate it indirectly using a narrow-linewidth 674 nm laser to address various electric quadrupole (E2) transitions to the meta-stable $D_{5/2}$ states in anticipation of future multi-ion experiments (and as we already rely on that 674 nm laser for state-selective readout). Hence, access to both a $|\Delta m| = 1$ and $|\Delta m| \in \{0, 2\}$ transition is required, preferably with a single laser beam for simplicity. Considering the E2 geometry factor while minimising the magnetic field sensitivity, we usually use $|D_{5/2}, m_J = -3/2\rangle$ as the shared upper state, addressing transitions from the ground state to it using a laser beam angled perpendicular to the magnetic field and linearly polarised. A π pulse on the $|S_{1/2}, m_J = 1/2\rangle \to |D_{5/2}, m_J = -3/2\rangle$ transition maps the Zeeman qubit state onto the less magnetically-sensitive $|S_{1/2}, m_J = -1/2\rangle \leftrightarrow |D_{5/2}, m_J = -3/2\rangle$ qubit, where it can then be further manipulated. The angle of polarisation tunes the relative strength of the two transition classes (and is usually set to favour $|\Delta m| = 1$ transitions for multi-qubit gates).

The 674 nm laser source is a DPSS-pumped Ti:Sapphire laser[9], stabilised to an "ultra-low expansion" (ULE) glass cavity[10] using the Pound–Drever–Hall technique augmented by an additional frequency offset tone, with feedback applied via an external AOM (for fast actuation) and to piezoelectric actuators controlling the laser cavity length. The frequency of the offset tone, and hence the frequency shift between locked laser and cavity resonance, is slowly ramped to compensate a slowing frequency drift of the reference cavity presumably from relaxation of manufacturing stress (0.199 mHz/s in September 2017, 0.055 mHz/s in December 2021). The low-bandwidth feedback to the laser cavity length is implemented using a custom digital feedback controller[11] that also monitors the high-finesse cavity transmission to remedy occasional loss-of-lock events (likely in part caused

---

[8]There also happens to be a 729 nm ($S_{1/2} \leftrightarrow D_{5/2}$) clock transition in ^{43}Ca$^+$ close to this field [Ben+07]; in a future expansion, the field could be re-tuned to the exact point of vanishing first-order field sensitivity at 0.496 mT with minimal changes.

[9]Supplied by a Scottish manufacturer who shall remain nameless, as I would otherwise feel obliged to comment on the abysmal reliability record of this laser system.

[10]Cavity and system integration provided by Stable Laser Systems, Boulder, CO, United States.

[11]The controller is based on the *Stabilizer* microcontroller card from the Sinara open-hardware ecosystem, extended with isolated digital outputs and executing custom firmware. The digital outputs are used to control a fast analog controller for the main AOM-based lock and provide lock status signals to the ion trap nodes. Automatic lock re-acquisition is coordinated by a state machine executed on a desktop PC, as selecting the correct target mode relies on wavemeter measurements and can require digital adjustments to the Ti:Sapphire laser mode selection etalon.

Lower qubit state	Transition magnetic field sensitivity	Coherence time		Error per Clifford gate	
		Alice	Bob	Alice	Bob
$S_{1/2}, m_J = -1/2$	$-11.2\,\mathrm{MHz/mT}$	$7.5\,\mathrm{ms}$	$4.2\,\mathrm{ms}$	5.9×10^{-4}	4.9×10^{-4}
$S_{1/2}, m_J = +1/2$	$-39.2\,\mathrm{MHz/mT}$	$1.7\,\mathrm{ms}$	$1.0\,\mathrm{ms}$	7.5×10^{-4}	10×10^{-4}

Table 3.1: Measured qubit performance on 674 nm transitions with $D_{5/2}, m_J = -3/2$ as the upper qubit state. The $1/e$ coherence times were extracted from a Gaussian fit to the contrast decay in a Ramsey experiment with varying wait duration, without any dynamical decoupling pulses. The gate errors are obtained via randomised benchmarking [Kni+08]; in the chosen decomposition of the single-qubit Clifford group, the mean physical pulse area across all elements is 0.58π.

by acoustic disturbances[12]); a real-time signal notifies the ion-trap nodes of this so that data acquisition can be momentarily paused.

To stabilise the magnetic fields at the ion location, we use high-precision feedback controllers acting onto the quantisation coil currents (including active feed-forward to compensate for ambient fields generated by 50 Hz power mains)[13].

For manipulation of the $^{43}\mathrm{Ca}^+$ qubit state, including $^{43}\mathrm{Ca}^+$–$^{88}\mathrm{Sr}^+$ mixed-species entangling gates, two beams at $\approx 402\,\mathrm{nm}$ and spaced by 3.2 GHz are generated from a single external-cavity diode laser at $\approx 804\,\mathrm{nm}$ (separately amplified and frequency-doubled after applying the frequency splitting) to drive Raman transitions and induce spin-dependent light shifts; their use is not discussed further here.

3.2.4   Single-photon generation and analysis

Ion-photon entanglement generation (see chapter 4) starts with preparation of the ion in $\left|5s\,^2S_{1/2}, m_J = -\frac{1}{2}\right\rangle$ the excitation of the $^{88}\mathrm{Sr}^+$ ion to the short-lived $\left|5p\,^2P_{1/2}, m_J = \frac{1}{2}\right\rangle$ state. This is achieved via a $\sim 5\,\mathrm{ps}$, σ^+-polarised laser pulse resonant with the 422 nm transition, generated using a mode-locked Ti:Sapphire laser with a custom pulse-picking and frequency doubling system [Nad16]. Non-linear effects (self-phase modulation) in the regular step-index fused-silica single-mode fibres used to deliver the beams to the trap complicate the ion dynamics; we typically work with excitation probability well in excess of 95 %.

Photons emitted into free space during spontaneous decay, entangled in polarisation with the resulting $S_{1/2}$ Zeeman state according to the dipole selection rules, are collected by custom-designed lens objectives[14] (numerical aperture 0.6) perpendicular to the magnetic

---

[12]As supplied by Stable Laser Systems, only the reference cavity is contained in an acoustically shielded enclosure, but not the laser head. Measurements with noise played back from a loudspeaker revealed a region of high acoustic susceptibility, including several resonances, around 1–5 kHz.

[13]The controllers are similar to that described in ref. [Mer+19], but implemented as an add-on circuit board to the Sinara *Stabilizer* open-hardware controller platform. To apply the power-line harmonics compensation, we digitally modulate the controller set point.

[14]Designed and manufactured by Photon Gear Inc., Ontario, NY, United States.
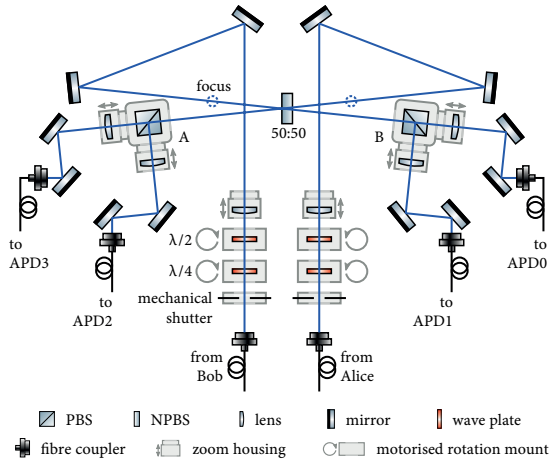
Figure 3.7: Beam paths of the central heralding station. A shallow angle of incidence on the non-polarising 50:50 beamsplitter minimises polarisation-dependent effects. The single-photon detectors APD0 and APD3 measure horizontally polarised photons (in the diagram plane), detectors APD1 and APD2 vertically polarised photons. (Figure adapted from ref. [Ste19].)

	Alice	Bob
$\lambda/4$	0.217(1)	0.241(1)
$\lambda/2$	0.449(1)	0.471(1)

Figure 3.8: In-situ measurement of the retardances of the input wave plates, given in units of 2π, as relevant for ion–photon state tomography [Ste19].

pol.	R	T
H	0.4693	0.5307
V	0.4717	0.5283

Figure 3.9: Measured transmittance and reflectivity of the non-polarising plate beamsplitter [Ste19].

pol.	PBS A	PBS B
H	12 500:1	3000:1
V	700:1	1900:1

Figure 3.10: Measured extinction ratios of the polarising beam-splitter cubes [Ste19].

field. They provide near-diffraction-limited performance at an input-side numerical aperture of 0.6 (specified: ≤ 0.08 waves rms at 403.1 nm). The image-side numerical aperture is 0.09, designed to match the ion emission into standard, commercially available step-index fused-silica fibres.

The objective lens is mounted on a five-axis stage for fine control over translational and rotational alignment. After fitting the observed point spread function with a low-order Zernike polynomial model similar to ref. [Won+16], a cylindrical lens is inserted near the image plane to correct for residual aberrations likely caused by varying thickness across the vacuum viewport (see refs. [Ste19; Nic22] for a detailed description).

The tip of the collection fibre is mounted on a three-axis positioning stage, which can be adjusted using open-loop piezoelectric actuators to correct for mechanical relaxation and slow, thermal drifts. Automated calibration experiments that periodically track the optimum in fluorescence coupled to the Bell-state analyser detectors allows us to retain good fibre-coupling efficiency over many days without manual intervention.

The trap nodes are connected to a central heralding station (or: "Bell-state analyser", "entangler") using 2×1.75 m fused silica single-mode fibres. There, a 50 : 50 beamsplitter and two polarising beamsplitters implement a partial Bell-basis measurement (see fig. 3.7). Real optical components are slightly (some significantly) non-ideal; measurements of the used wave plates and beamsplitters are given in figs. 3.8 to 3.10. At the output, the four beams are coupled into short single-mode fibres before being focussed onto four avalanche photodiode (APD) single-photon detectors connected to heralding logic implemented in

the node-local FPGA control systems. This significantly attenuates the impact of any mode mismatch at the central 50 : 50 beamsplitter on the remote entanglement fidelity (at the cost of a ~ 10 % decrease in single-photon efficiency through re-coupling losses). A pair of quarter- and half-wave plates compensates stress birefringence in the input fibres (and is used to select a photon measurement basis for experiments probing the ion–photon entanglement); their orientation is calibrated daily to track slow mechanical/thermal drifts.

The peak overall single-photon efficiencies observed (from picosecond-pulse excitation to detecting a click on one of the APDs) were $\approx 2.4\%$ in either system. Known contributions to this are state preparation and excitation ($\approx 97\%$), decays to $D_{3/2}$ instead of $S_{1/2}$ (95 %), overlap of the ion dipole image with the Gaussian fibre mode (73 %), bulk transmission through the anti-reflection-coated single-mode fibres ($\approx 98.7\%$), transmission and re-coupling of the heralding station ($\approx 90\%$), and quantum efficiency of the used APD detectors ($\approx 69\%$). A factor $\approx 58\%$ between experiment and theory remains, presumably mostly due to imperfect alignment and residual aberration in the high-NA optics reducing the ion-to-fibre coupling efficiency. Most of the data presented in this thesis was acquired at significantly lower efficiencies, however, due to miscellaneous alignment drifts across several months since the last careful alignment; in particular, the collection optics in Bob had accidentally been misaligned quite substantially.

3.2.5 Real-time control

The experimental hardware is controlled by a customised version of the ARTIQ ("Advanced Real-Time Infrastructure for Quantum Physics") open-source control system. High-level management and scheduling functions are performed by a PC, which also controls a large number of peripherals across the local network; the real-time infrastructure is based on several field-programmable gate array (FPGA) controller boards connected by fibre-optic links, which also contain an embedded processor that executes experiment-specific programs (just-in-time-compiled from a custom dialect of Python) close in latency to the FPGA event distribution network and hardware interface blocks. Most hardware is openly designed, engineered and documented as part of the ARTIQ/Sinara open hardware ecosystem, and was developed in collaboration with the Warsaw University of Technology and other institutions (in part during this project).

At 1 MHz entanglement generation attempt rate, latencies in AOMs, optical fibres, cabling, but also the digital control layer are no longer negligible. During the excitation cycle (see e.g. fig. 6.2), a custom sequencer core in the main node control FPGAs directly takes over the relevant logic signals (AOM RF switches, APD inputs) from the ARTIQ event distribution layer to execute a hard-coded state machine with minimal branching latency overhead. Furthermore, the 80 MHz cycle of the mode-locked 422 nm laser is not directly synchronised with the control system for technical simplicity, leading to an effective jitter of 12.5 ns in the timing of the excitation pulse relative to the other real-time hardware

(which is synchronised to $\ll 1$ ns). The custom FPGA core thus references all coincidence detection windows to an auxiliary photodiode signal giving the timing of the excitation pulse. The sequencer cores on each node establish real-time synchronisation with each other across a few digital signal lines[15].

A classical network link (Ethernet) between Alice's and Bob's master PCs is used for high-level coordination, as well as to exchange the messages required for the DIQKD post-processing steps in chapter 7. We have developed custom coordination and scheduling logic to, for instance, handle interruptions due to ion-loss events or periodic recalibrations (some of which need access to shared resources, e.g. the heralding station for fibre alignment) when running two-node experiments. Except for rare failures, the system is capable of autonomous operation for days at a time[16]; a suite of automated calibrations and diagnostics tests is executed twice every night.

---

[15]Were the nodes to be deployed in distant locations, all real-time signalling would sensibly (and easily) be multiplexed onto a single optical fibre link connecting the nodes.

[16]The capability for stand-alone operation has thus far been hard to rigorously quantify, as various parts of the system are actively being worked on still, which tends to introduce new issues.

4 Theory of remote ion–ion entanglement through spontaneous emission

In this chapter, we discuss the generation of entanglement between the state of an atom and the polarisation of a photon emitted during spontaneous decay on a transition between two levels with total angular momentum quantum number $1/2$ – here, the 422 nm $P_{1/2} \rightarrow S_{1/2}$ transition in $^{88}\mathrm{Sr}^+$ –, and the establishment of entanglement between two such atoms through subsequent measurement of the photons (entanglement swapping).

In particular, we will consider the case where the atomic emission is collected by a lens from a cone perpendicular to the ambient magnetic field. In prior work, it had already been known that coupling the emission into a single-mode optical fibre suppresses unwanted contributions for spontaneous emission from atoms with nuclear spin $1/2$ [Luo+09; KMK11]. Somewhat unexpectedly, the same mechanism, based on the radial symmetry of the optical fibre modes, also suppresses errors from polarisation mixing and imbalance amplitude imbalances in our case; we show that collection from an arbitrarily large solid angle incurs, in principle, no loss in entanglement fidelity (§4.2).

As any limitations in fidelity are thus predicted to be technical in origin, we are motivated to consider quantitatively the effect of an extensive list of experimental imperfections (§4.3). In this high-fidelity regime, it is also no longer useful to consider the Bell-basis measurement on the photon polarisations which converts the ion–photon entanglement to entanglement between the remote ions to be perfect. We introduce a generic formalism for the treatment of the heralding station (§4.4), and similarly catalogue the effects of a number of component imperfections or residual distinguishability between the input photons (§4.5).

We will analyse the system as a concatenation of independent steps: spontaneous emission into vacuum, the collection of a fraction of the emission into an optical single-mode fibre, and then the interference of such wavepackets from two network nodes, entangled with the ion in each trap, in the linear-optics partial Bell-basis measurement apparatus. Such an approach is possible in this case as the photonic mode structure at the ion position is not appreciably modified by the presence of the non-reflective collection optics (as compared to, for instance, an ion in an optical cavity), and there is no driven action or fast back-action on the atomic system during the emission of the photon wavepacket. There are
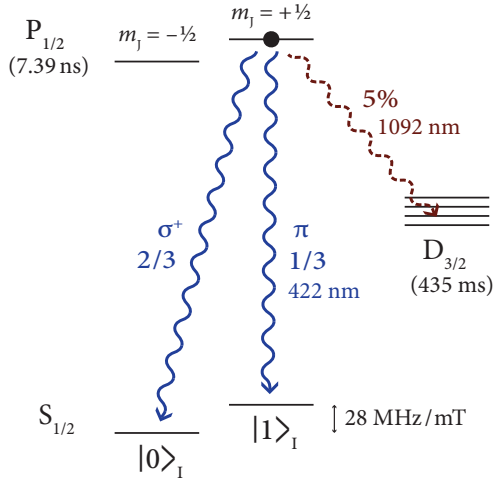
Figure 4.1: Diagram of the low-lying energy levels in $^{88}Sr^+$, as relevant to ion–photon entanglement generation.
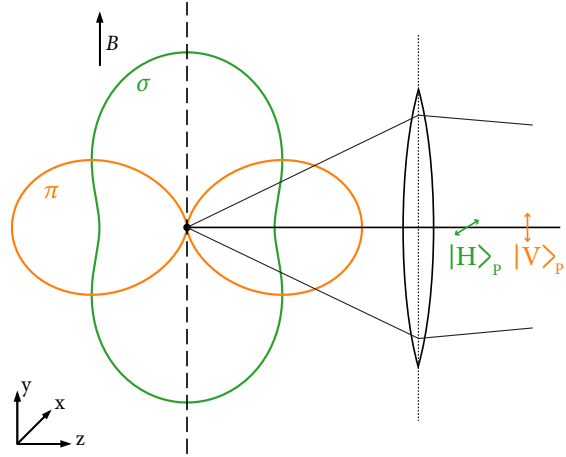
Figure 4.2: Illustration of the ion–photon entanglement mechanism: The intensity of π electric dipole transitions is twice that of σ transitions perpendicular to the magnetic field axis (compensating the difference in Clebsch–Gordan coefficients), and the polarisations are orthogonal.

two main advantages to this: First, this approach is conceptually and notationally simple; the spatial aspect is treated solely when discussing the ion–photon fibre coupling, the temporal complexity when later describing the measurements. Secondly, this also allows us to easily "plug in" experimental data, that is, model the action of the photonic measurement setup on particular ion–photon states described by density matrices obtained through state tomography. For extensions beyond these assumptions, the photon mode (overlap) considerations described here could still form the basis of a more tightly integrated analysis, e.g. using a quantum jump[1] approach (see e.g. refs. [PK98; WWM02; MB08]).

4.1   Spontaneous emission of 422 nm photons in $^{88}Sr^+$

We consider the particular case relevant for our experiment, that of fig. 4.1, where a $^{88}Sr^+$ ion is excited to the state $|e\rangle = \left|5p\,^2P_{1/2}, m_J = \frac{1}{2}\right\rangle$. There is a strong electric dipole (E1) coupling connecting $|e\rangle$ and the two Zeeman-split ground qubit states $|0\rangle = \left|5s\,^2S_{1/2}, m_J = -\frac{1}{2}\right\rangle$ and $|1\rangle = \left|5s\,^2S_{1/2}, m_J = \frac{1}{2}\right\rangle$, so the ion will decay to the ground state a short time later under the emission of a 422 nm photon.

The E1 coupling of $|e\rangle$ to states in the $D_{3/2}$ level is non-negligible as well, but as these states are metastable and the 1092 nm transition wavelength is vastly different from the dominant $S_{1/2}$ decay channel (to where such photons will not be collected and detected in our apparatus at all), we focus entirely on the 422 nm decays here. Defining the branching

---

[1] The technique of deriving the dynamics of open quantum systems explicitly from unitary dynamics interrupted by quantum jumps is variously also known as the "Monte–Carlo wave function" or "quantum trajectories" approach.

fraction $\eta_{422} := \sum_{i \in S_{1/2}} A_{ei}/\Gamma$ (in obvious notation for the sum over final states), the $1 - \eta_{422} = 5\,\%$ of decays to one of the $D_{3/2}$ states only contribute a small additional loss to the overall efficiency of protocols heralded on the detection of a 422 nm photon[2].

In the chosen naming convention, $|0\rangle$ corresponds to the decay on the σ^+ transition, $|1\rangle$ on the π transition. The atomic dipoles from eq. (2.64) are

$$\boldsymbol{d}_{\sigma^+} := \boldsymbol{d}_{0e} = e\langle 5s\,^2S_{1/2}\|T^{(1)}\|4p\,^2P_{1/2}\rangle C_0 \hat{\boldsymbol{e}}_{+1}, \quad C_0 = -\sqrt{\frac{2}{3}},$$
$$\boldsymbol{d}_{\pi} := \boldsymbol{d}_{1e} = e\langle 5s\,^2S_{1/2}\|T^{(1)}\|4p\,^2P_{1/2}\rangle C_1 \hat{\boldsymbol{e}}_0, \quad C_1 = \sqrt{\frac{1}{3}}. \tag{4.1}$$

From the ratio of (squared) Clebsch–Gordan coefficients, $|C_0/C_1|^2 = 2$, σ^+ decays are twice as likely as π decays. Considering long times $t \gg \Gamma^{-1}$ such that the remaining excited state population is negligible, and tracing out the photon modes, the ion state is thus

$$\rho_I \approx \eta_{422}\left(\frac{2}{3}\,|0\rangle\langle 0| + \frac{1}{3}\,|1\rangle\langle 1|\right) + (1 - \eta_{422})\,\rho_{1092,I}, \tag{4.2}$$

where we have elided the result of the $D_{3/2}$ decays as $\rho_{1092,I}$. As we shall see in detail in §4.2, the difference in the Clebsch–Gordan coefficient magnitude will be exactly balanced out by the dipole emission patterns for observation perpendicular to the magnetic field, though.

The notation for the overall state, in particular the photonic modes, can be refined to more clearly elucidate some key aspects. First, we will – for the entire rest of the chapter –, switch into an interaction picture with respect to the internal ion state by eliding the phase factors from eq. (2.78). This is equivalent to applying all later operations on the Zeeman qubit in a rotating frame *anchored to the laboratory frame at the time of the excitation pulse*[3], $t = 0$. Furthermore, as the frequency difference $|\omega_{e0} - \omega_{e1}| = |\omega_0 - \omega_1|$ between the two 422 nm decays (given by the Zeeman splitting) is very small compared to the absolute optical frequency, we will factor out a common reference frequency ω_n (and corresponding angular wavenumber k_n), which for later convenience we take to be the transition frequency at zero static magnetic field B, such that $\omega_{ei} = \omega_n + \delta_i = \omega_n + \nu_{ei}B$,

---

[2]In "single-photon herald" remote entanglement schemes, where only one 422 nm photon would be collected from two ion trap nodes, it might be advantageous to exclude such 1092 nm decays through an extra protocol success verification step, where undesired decays cause the protocol to start again in the same way as when no herald occurs. This could be achieved through shelving of both qubit states to the $D_{5/2}$ manifold followed by $S_{1/2} \leftrightarrow P_{1/2} \leftrightarrow D_{3/2}$ fluorescence detection.

[3]As we will see, the excitation pulse is not actually a particular meaningful timing reference in practice, as the ion–photon state phase will turn out to depend on the photon detection time, and the ion–ion remote entangled state will have no phase evolution at all for matched qubit frequencies. It is nevertheless important not to lose track of this for conceptual clarity.

where ν_{ei} is the overall magnetic field sensitivity of the respective transition. We then split up the mode functions from eq. (2.82) according to

$$
f_i(\boldsymbol{r}, t) \approx \underbrace{\sqrt{\pi}\sqrt{\frac{\hbar \omega_n}{\varepsilon_0 c}}}_{=: \mathscr{E}} \phi_i\left(t - \frac{r}{c}\right) \boldsymbol{\varepsilon}_i(\boldsymbol{r}) e^{-i\omega_n t}
$$

$$
\boldsymbol{\varepsilon}_i(\boldsymbol{r}) := \sqrt{\frac{3}{8\pi}}\left(\hat{\boldsymbol{d}}_{ie} - (\hat{\boldsymbol{d}}_{ie} \cdot \hat{\boldsymbol{r}})\hat{\boldsymbol{r}}\right)\frac{e^{ik_n r}}{r} \tag{4.3}
$$

$$
\phi_i(t) := \sqrt{\Gamma} e^{-i\delta_i t}\, e^{-\frac{\Gamma}{2}t}\, \Theta(t),
$$

where we have chosen the normalisation of the time-independent spatial mode $\boldsymbol{\varepsilon}_i(\boldsymbol{r})$ such that the "flux" $|\boldsymbol{\varepsilon}_i|^2$ integrated over spherical shells is 1, that of the temporal envelope $\phi_i(t)$ such that $\|\phi_i\|^2 = \langle\phi_i|\phi_i\rangle = \int_{-\infty}^{\infty}|\phi_i(t)|^2 \mathrm{d}t = 1$, and have made the approximation $\omega_{ei} \approx \omega_n$ again for the prefactor for purely æsthetic reasons: to highlight that the only appreciable contribution of the small difference in transition frequencies is a modulation of the phase of the temporal envelope. Because the modes are normalised, the absolute dipole strength of the dipole moment for the particular transition does not enter the expression (except through its contribution to the total excited state linewidth Γ).

Considering only the two decay channels to the $S_{1/2}$ ground state, we can write a normalised state[4]

$$
|\psi^{422}\rangle = \sqrt{\frac{2}{3}}\,|0\rangle \otimes (\tilde{a}_0^\dagger\,|0\rangle) + \sqrt{\frac{1}{3}}\,|1\rangle \otimes (\tilde{a}_1^\dagger\,|0\rangle), \tag{4.4}
$$

such that, neglecting the coherences between the $S_{1/2}$ and $D_{3/2}$ states[5], the overall state is

$$
\rho_{IP} = \eta_{422}\,|\psi^{422}\rangle\langle\psi^{422}| + (1 - \eta_{422})\,\rho_{IP}^{1092}, \tag{4.5}
$$

where we again shall not concern ourselves any further with the population ρ_{IP}^{1092} in the $D_{3/2}$ states.

To discuss the state of the system in an intuitive, yet formally precise fashion, we briefly transition to the photon wavefunction formalism, where the value of the photonic wavefunction is given by $\boldsymbol{E} + ic\boldsymbol{B}$. We can write the electric part $|\psi_E(\boldsymbol{r}, t)\rangle$ of the state (4.4) as

$$
|\psi_E^{422}(\boldsymbol{r}, t)\rangle = \mathscr{E}e^{-i\omega_n t}\left(\sqrt{\frac{2}{3}}\phi_0(t - r/c)\,|0\rangle \otimes \boldsymbol{\varepsilon}_{\sigma^+}(\boldsymbol{r}) + \sqrt{\frac{1}{3}}\phi_1(t - r/c)\,|1\rangle \otimes \boldsymbol{\varepsilon}_{\pi}(\boldsymbol{r})\right), \tag{4.6}
$$

such that, restoring the magnetic part, $|\psi^{422}(\boldsymbol{r}, t)\rangle = |\psi_E^{422}(\boldsymbol{r}, t)\rangle + i\hat{\boldsymbol{r}} \times |\psi_E^{422}(\boldsymbol{r}, t)\rangle$. If we take the static magnetic field to be oriented along $\hat{\boldsymbol{z}}$ as per the usual atomic physics

---

[4]Recall that we have switched into the interaction picture regarding the internal state of the ion.

[5]Not only do we ignore the $D_{3/2}$ states in what follows anyway; observing the phase between the levels would require absolute phase stability between the 687 nm laser that would presumably be used to address the $S_{1/2} \leftrightarrow D_{3/2}$ transition and the 422 nm picosecond pulsed laser, which would be a major technical hurdle.

convention, and consider the wave function along a line $\hat{r} = y\hat{y}$ perpendicular to the magnetic field, we have $\boldsymbol{\varepsilon}_{\sigma^+}(y\hat{y}) = -\sqrt{\frac{3}{16\pi}}\frac{e^{ik_n y}}{y}\hat{x}$ and $\boldsymbol{\varepsilon}_{\pi}(y\hat{y}) = \sqrt{\frac{3}{8\pi}}\frac{e^{ik_n y}}{y}\hat{z}$, and

$$|\psi_E^{422}(y\hat{y},t)\rangle = \frac{\mathscr{E}}{\sqrt{8\pi}}\frac{e^{i(k_n y - \omega_n t)}}{y}\left(-\phi_0\left(t - \frac{y}{c}\right)|0\rangle \otimes \hat{x} + \phi_1\left(t - \frac{y}{c}\right)|1\rangle \otimes \hat{z}\right). \qquad (4.7)$$

In other words, perpendicular to the magnetic field, the two decay channels map to orthogonal polarisations, and as $\phi_1(t) = e^{i(\delta_0 - \delta_1)t}\phi_0(t)$, occur equally often. Thus, we might consider the ion state to be maximally entangled with the photon polarisation, and suggestively write an "ion–photon state" like

$$|\Psi^{422}\rangle^{r\|\hat{y}} \sim \sqrt{\frac{1}{2}}\left(|0\rangle|H\rangle + |1\rangle|V\rangle\right). \qquad (4.8)$$

This does provide a good heuristic picture of the entanglement generation mechanism, but is only accurate in a very approximate sense. First, the relation only holds exactly orthogonal to the magnetic field, which is an infinitesimally small solid angle. Secondly, note that we already transitioned into the rotating frame of the internal ion states; the term $e^{i(\delta_0 - \delta_1)(t - r/c)}$ is an additional amount of differential phase accumulation that we need to actively correct if we want to avoid the state to dephase across different detection times t – or, equivalently, reference the rotating frame used for manipulations on the $\{|0\rangle, |1\rangle\}$ qubit to the laboratory rest frame[6].

4.1.1 Limitations of this treatment

At the very beginning of the derivation in §2.4.5, we assumed that the ion states out in the pure state $|e\rangle$, that is, neglected the excitation dynamics. The lifetime $1/\Gamma$ of the excited state is necessarily short (here, 7.39 ns for the $P_{1/2}$ level in ^{88}Sr⁺) as we wish to quickly extract spontaneous emission photons. As our excitation pulses are yet shorter, on the order of 5 ps in duration, neglecting the excitation dynamics like this is an excellent assumption. This is not necessarily the case for experiments that use nanosecond-scale pulses from a gated continuous-wave laser; in that case, the dynamics are more complex, but still relatively straightforward to analyse [Wan+11; Mül+17].

By writing $|e\rangle$ as the initial state without a phase prefactor, we have also chosen a phase reference for the emitted photon wavepackets (as the dynamics generated by eq. (2.66) are unitary, an additional phase would just carry through to the absolute phase of the emitted photon wavepackets). As we cannot observe an offset to the absolute optical phase

---

[6]Without the further manipulation and detection of the photon in the picture, this is necessarily a somewhat perfunctory description. It does turn out to be precise in the sense that after post-selecting experiments on the measurement of a photon of a given polarisation, the ion qubit state, in the rotating frame given by the photon detection time, is exactly as would have resulted from a maximally entangled state between the ion Zeeman qubit and an abstract, stationary polarisation qubit.

at 422 nm in our experiment (see §4.5.1), we are free to make this identification. In cases where the absolute phase does set the phase of the generated entangled state (e.g. for single-photon herald schemes), this means that the entangled state phase is referenced to the optical phase of the picosecond laser pulse at the ion which coherently excites the dissipatively prepared state $|0\rangle$ to $|e\rangle$.

Finally, note that we have neglected the centre-of-mass motion of the ion in this treatment. As mentioned in §2.4, if quantisation is carried out in a self-consistent fashion for the combined system, extra Röntgen-type terms coupling the atomic dipole to the magnetic field appear in the Hamiltonian even in the dipole approximation, but they are typically weak. The main addition to the picture is a momentum kick to the atom entangled with the photon propagation direction, but this can be incorporated without many additional complications (cf. [Bab84; RZ92; Lem+93; BBL93; Wil94; Fed+05], but note that these references discuss the situation for a free atom, while the position and momentum operators would here decompose into that of the quantised secular modes of motion).

4.2 Fibre-coupling of ion dipole emission

To calculate the focussed radiation fields corresponding to atomic dipole emission modes is a surprisingly rich problem. Many of the approaches that are widely used to describe optical systems apply only to the source-free case – that is, the homogeneous Maxwell equations –, and then provide approximate solutions by making additional assumptions, such as those of scalar fields and small-angle paraxial optics. In the present case, these clearly do not apply: The central region does contain a source (the atom), and we are interested in collecting as large a fraction of the approximately isotropic emission as is technically practical. As our aim is to study high-fidelity entanglement between atom and photon polarisation, we also do not want to neglect any polarisation effects.

In the following, we derive the fields close to the image of a lens in the full, vectorial picture. For this, we start from the spatial part of the atomic emission modes ϵ (eq. (4.3)), which coincide with the far-field radiation of a classical dipole emitter in vacuum. We then derive the action of an ideal optical system from geometric considerations. A brief look at the resulting fields in a collimated plane wave picture will already allow us to illustrate most effects of interest here. Finally, the same geometric considerations will allow us to propagate the field through a finite-conjugates imaging system in the complete vectorial formalism, including diffraction at the exit pupil of the system.

The resulting solution describes the field close to the focus under only mild assumptions (mostly the aperture size and the distance of object and image points to the optical system both being much larger than the wavelength), which we can then use to calculate the system performance over a large range of parameters, including after coupling into a single-mode optical fibre.
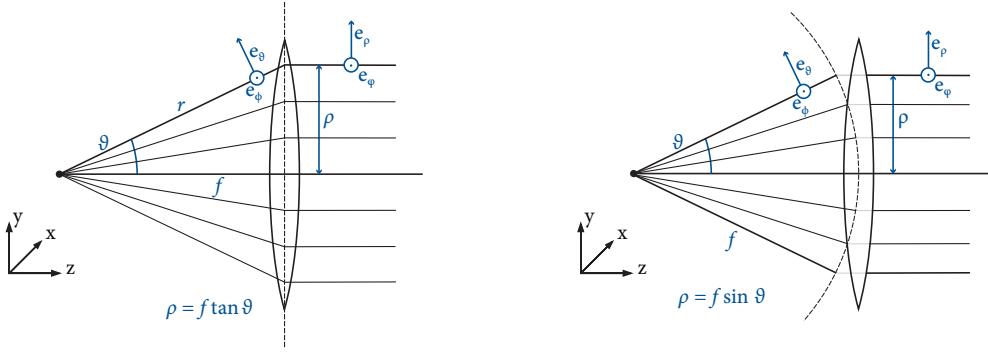
Figure 4.3: Two models for a geometrical optics treatment of point-source radiation collimated by an ideal lens. *Left:* An ideal thin lens. *Right:* An ideal lens following the Abbe sine criterion. Shown are also the unit vectors of the spherical/polar coordinate systems used in the text.

4.2.1   Ideal lenses

An ideal lens of focal length f, positioned at distance f from the ion, will collimate radiation emitted into the solid angle it subtends. The size of this collection cone is customarily specified in terms of the numerical aperture (NA), $n \sin \vartheta_0$, where ϑ_0 is the cone half-angle and n is the index of refraction of the medium (here, $n = 1$ throughout).

There are two possible ways of modelling such an ideal element, as shown in fig. 4.3: In both cases, the lens collimates all rays originating from the ion's position and exactly flattens out the spherical wavefronts. To describe this action, it is convenient to describe the fields in spherical coordinates (r, ϑ, ϕ), where $\vartheta = 0$ is aligned in the direction of the optical axis $\hat{z}$, and the field ϵ_B immediately after the lens is described in polar coordinates (ρ, φ), as shown in fig. 4.3. The spherical wavefront e^{ikr} is then taken to a phase uniform across the cross-section of the beam, which we can set to zero by choosing an appropriate reference plane.

Where the models differ is in the precise mapping of the coordinates. On the left of fig. 4.3, the result of ray tracing across an ideal, infinitely thin lens is shown, where the coordinates are transformed as

$$\rho = f \tan(\vartheta) \qquad \varphi = \phi. \tag{4.9}$$

On the right, a lens following the Abbe sine criterion is shown, which corresponds to refraction on a reference sphere around the focal point, and gives another possible coordinate transformation as

$$\rho = f \sin(\vartheta) \qquad \varphi = \phi. \tag{4.10}$$

The coordinate vectors then simply transform as $\hat{e}_\vartheta \mapsto \hat{e}_\rho$, $\hat{e}_\phi \mapsto \hat{e}_\varphi$ in both cases.

To determine how the field magnitudes are transformed, consider the flux $|\epsilon|^2 r^2 \sin \vartheta \, d\vartheta \, d\phi$ through an infinitesimal region on a spherical shell of radius r in the ion emission, which

– by energy conservation – must be equal to the flux through the corresponding cross-section of the collimated beam, $|\epsilon_B|^2 \, \rho \, d\rho \, d\varphi$. Hence, for eq. (4.9)

$$\frac{|r\,\epsilon(r,\vartheta,\phi)|^2}{|\epsilon_B(\rho,\varphi)|^2} = \frac{\rho}{\sin\vartheta}\frac{d\rho}{d\vartheta} = \frac{f^2}{\cos^3\vartheta}, \tag{4.11}$$

while for the sine condition from eq. (4.10),

$$\frac{|r\,\epsilon(r,\vartheta,\phi)|^2}{|\epsilon_B(\rho,\varphi)|^2} = \frac{\rho}{\sin\vartheta}\frac{d\rho}{d\vartheta} = f^2\cos\vartheta. \tag{4.12}$$

The respective scaling factors for the amplitudes are taken to be real and positive by convention, $(\cos\vartheta)^{3/2}$ and $\sqrt{1/\cos\vartheta}$, and are sometimes referred to as *apodisation factors*.

Both approaches have been used to model imaging of point-like sources (cf. [SG93]). While the sine criterion can be derived from the requirement of stigmatic imaging of off-axis points, here only a single point-like source is relevant, and the two models coincide in the limit of small collection angles. In the absence of a particular physical imaging system to model, either is a sensible choice.

While the fields ϵ_B are only defined in the plane where the lens acts, it is nevertheless instructive to compute them and inspect the results. Note that eq. (4.9) maps the field to the whole plane $\rho \in [0,\infty)$, whereas eq. (4.10) only maps into the disc $\rho \in [0,r]$, in accordance to the geometrical construction from fig. 4.3. Correspondingly, the apodisation factor for the latter diverges at the boundary $\vartheta = \frac{\pi}{2}$. Because of this, and the slightly more straightforward geometrical interpretation of the former, we will assume the tangent condition for now, before switching to the sine condition for the complete imaging treatment in the next section.

By expressing the atomic dipoles $\hat{\boldsymbol{d}}_{\sigma^+}$ and $\hat{\boldsymbol{d}}_{\pi}$ from eq. (2.15) using the coordinate vectors

$$\hat{\boldsymbol{e}}_{\vartheta} = \begin{pmatrix} \cos\vartheta\cos\phi \\ \cos\vartheta\sin\phi \\ -\sin\vartheta \end{pmatrix}, \qquad \hat{\boldsymbol{e}}_{\phi} = \begin{pmatrix} -\sin\phi \\ \cos\phi \\ 0 \end{pmatrix}, \tag{4.13}$$

and further substituting for dimensionless units $\rho = f\chi$, the collimated fields are obtained as[7]

$$\epsilon_{B,\sigma^+}(\chi,\varphi) = \sqrt{\frac{3}{16\pi}}\frac{1}{\sqrt{1+\chi^2}^{\frac{3}{2}}}\left(\frac{\cos\varphi - i\chi}{\sqrt{1+\chi^2}}\,\hat{\boldsymbol{e}}_{\chi} - \sin\varphi\,\hat{\boldsymbol{e}}_{\varphi} \right) \tag{4.14a}$$

$$\epsilon_{B,\pi}(\chi,\varphi) = \sqrt{\frac{3}{8\pi}}\frac{1}{\sqrt{1+\chi^2}^{\frac{3}{2}}}\left(\frac{\sin\varphi}{\sqrt{1+\chi^2}}\,\hat{\boldsymbol{e}}_{\chi} + \cos\varphi\,\hat{\boldsymbol{e}}_{\varphi} \right). \tag{4.14b}$$

---

[7]Strictly speaking, the coordinate vectors in eq. (4.14) are well-defined only for $\chi > 0$, but the fields can be trivially continued to the origin.
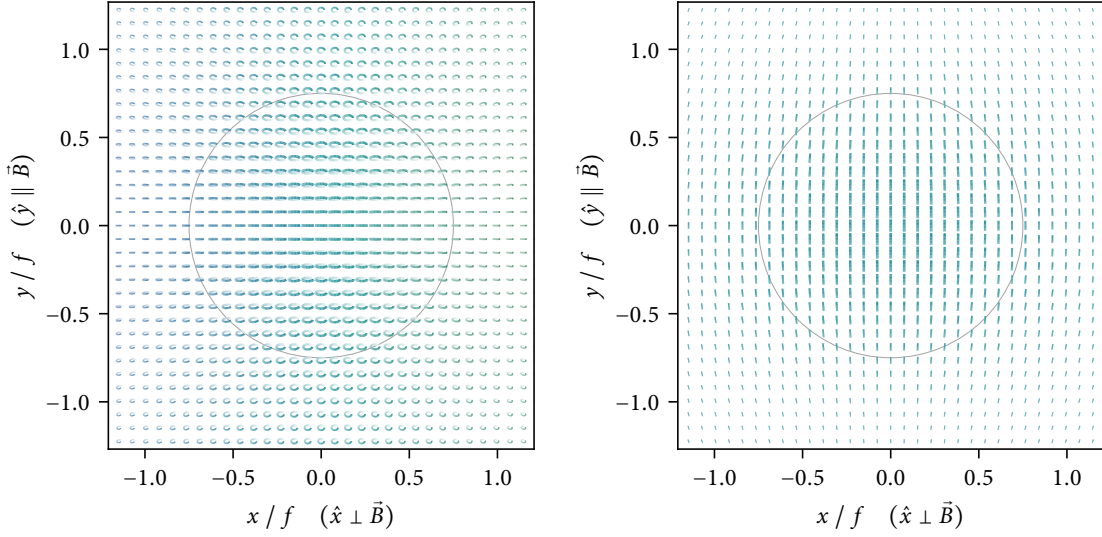
Figure 4.4: Ion emission pattern for σ^+ (left) and π (right) dipole transitions, after collimation by an ideal thin lens of focal length f (grey circles correspond to NA 0.6). The symbol shapes represent the electric field polarisation at each point in the plane, and their size and transparency are proportional to the amplitude. Outside the coordinate origin, the two modes are no longer perfectly orthogonal, but the unwanted components are symmetrical.

The resulting fields are shown in fig. 4.4. While the polarisations are indeed orthogonal in the origin, they now vary across the spatial mode. If the polarisation degree of freedom is later manipulated in free space without any spatial filtering, the ion–photon state can be described as an ensemble average over the collimated beam. If photons are collected from a cone of half-angle $\vartheta_0 > 0$, the (unnormalised) density matrix ρ_{IP} is

$$\rho_{IP} = \int\limits_0^{\tan\vartheta_0} \int\limits_0^{2\pi} |\psi(\chi,\varphi)\rangle\langle\psi(\chi,\varphi)| \, d\varphi \, \chi \, d\chi, \text{ where} \tag{4.15a}$$

$$|\psi(\chi,\varphi)\rangle = |0\rangle_I \, |\boldsymbol{\epsilon}_{B,\sigma^+}(\chi,\varphi)\rangle_P + |1\rangle_I \, |\boldsymbol{\epsilon}_{B,\pi}(\chi,\varphi)\rangle_P. \tag{4.15b}$$

The result is a non-maximally-entangled, mixed state; an example density matrix for numerical aperture $\sin\vartheta_0 = 0.6$ is shown in figure 4.5. The overall collection efficiency increases with ϑ_0, but so does the imbalance between decay channels and, crucially, the mixing of polarisations corresponding to each. This significantly decreases the resulting fidelities, shown quantitatively in fig. 4.6.

Note, however, that the unwanted components are highly symmetric, which is e.g. visible in fig. 4.4. In the next section, we will show that this observation carries through to the focus of an imaging system.

4.2.2 Derivation of general dipole image fields

Consider now the complete imaging system from fig. 4.7. We have again chosen the optical axis to be along $\hat{z}$. The dipole is imaged by an aplanatic optical system of focal length ratio

	$\langle 0H \vert$	$\langle 1H \vert$	$\langle 0V \vert$	$\langle 1V \vert$
$\vert 0H \rangle$	0.499			—
$\vert 1H \rangle$	$0.000 \cdot e^{i0.000\pi}$	0.001		
$\vert 0V \rangle$	$0.000 \cdot e^{i0.000\pi}$	$0.001 \cdot e^{i0.000\pi}$	0.047	
$\vert 1V \rangle$	$0.451 \cdot e^{i0.000\pi}$	$0.000 \cdot e^{i0.000\pi}$	$0.000 \cdot e^{i0.000\pi}$	0.452

0.5   0.25   0.1   0.01
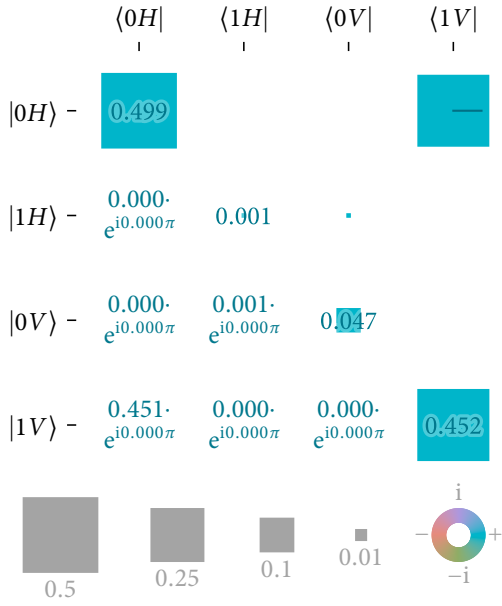
Figure 4.5: Density matrix ρ_{IP} for the ion–photon state in free space after an ideal NA 0.6 lens perpendicular to the magnetic field, conditional on the presence of a photon in the mode captured by the detector. There is a ∼ 10 % imbalance in decay channels, and the polarisation states corresponding to either decay channel are mixed (for σ^+ more so than for π).
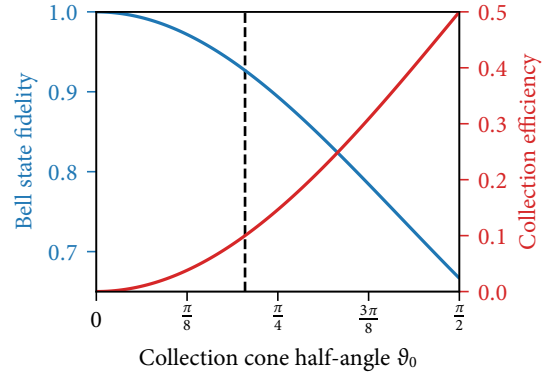
Figure 4.6: Fidelity to closest maximally entangled state (blue, left) and collection efficiency (red, right) when collecting ion emission in free space. The dashed line denotes NA 0.6, where the error has reached 6.5 %.

f'/f, now assumed to fulfil Abbe's sine condition[8], and to be free of aberrations.

To calculate the field in a point near the geometric image of the atomic dipole, we will follow the formalism developed by E. Wolf [Wol59], which he had applied (together with B. Richards [RW59]) to calculate the high-NA focus of a linearly polarised laser beam already in 1959. It has remained in wide-spread use for calculations involving high-NA imaging since; a modern overview can be found in chapters 3 and 4 of ref. [NH06]. In particular, Sheppard and Wilson [SW82] and later Enderlein [End00] have applied it to dipole emitters in calculations similar to the following, albeit mostly from a microscopy point of view, assuming in the first case uniform polarisation, and centering the discussion on the volumetric focussing properties for fluorescence correlation microscopy in the second.

The basic strategy is to propagate the object-space field to the image-side reference surface by tracing rays according to geometric optics, and then calculate the corresponding far field diffraction integral to calculate the field in the chosen image-space point.

We thus start as before with the dipole $\hat{\boldsymbol{d}}$ in the centre of the object-space reference sphere with spherical coordinates (r, ϑ, ϕ), $\vartheta \in [0, \pi/2)$, $\phi \in [0, 2\pi)$, along the optical axis

---

[8]We treat a system conforming to the sine condition here, as conventional designs derived from microscope objectives tend to be optimised in this way [GKG97]. Only the quantitative details of the image-plane intensity profile would be affected by a different choice, not its symmetries, so the differences are fairly inconsequential for the present discussion anyway.
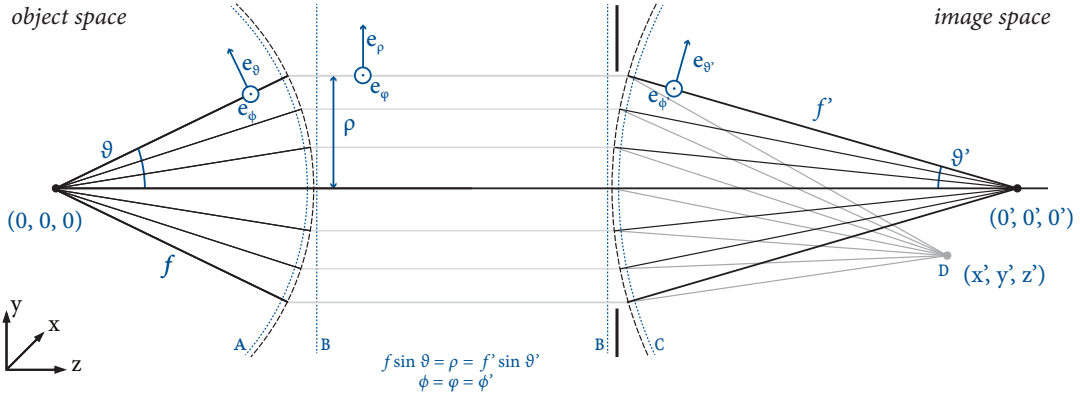
Figure 4.7: A source at object-space coordinates $(0,0,0)^T$ is imaged by an aberration-free (i.e., aplanatic) system, defining the coordinate origin of the image space. The field amplitude at a point $(x', y', z')^T$ in the image space can be calculated by geometrically propagating the spherical wave incident at the input reference sphere **A** to the output reference sphere **C**, and then evaluating a diffraction integral (see text).

$\hat{z}$. Evaluating eq. (4.3) gives the field just before the first reference sphere as

$$\boldsymbol{\epsilon}_A(\vartheta, \phi) = -\sqrt{\frac{3}{8\pi}} \frac{\mathrm{e}^{ikf}}{f} \left((\hat{\boldsymbol{e}}_\vartheta \cdot \hat{\boldsymbol{d}}) \hat{\boldsymbol{e}}_\vartheta + (\hat{\boldsymbol{e}}_\phi \cdot \hat{\boldsymbol{d}}) \hat{\boldsymbol{e}}_\phi \right), \tag{4.16}$$

where $\hat{\boldsymbol{e}}_\vartheta = \hat{\boldsymbol{e}}_\vartheta(\vartheta, \phi)$ and $\hat{\boldsymbol{e}}_\phi = \hat{\boldsymbol{e}}_\phi(\vartheta, \phi)$ are the respective coordinate vectors.

To derive the field just after the second reference sphere, we conceptually split the optical system into two parts, where the field is refracted at a reference sphere twice as described in the last section (in reverse the second time). Assuming now that the Abbe sine condition holds (with accordingly different apodisation), the field right after refraction at the surface is then

$$\boldsymbol{\epsilon}_B(\rho, \varphi) = -\sqrt{\frac{3}{8\pi}} \frac{\mathrm{e}^{ikf}}{f} \sqrt{\frac{1}{\cos\vartheta}} \left((\hat{\boldsymbol{e}}_\vartheta \cdot \hat{\boldsymbol{d}}) \hat{\boldsymbol{e}}_\vartheta + (\hat{\boldsymbol{e}}_\phi \cdot \hat{\boldsymbol{d}}) \hat{\boldsymbol{e}}_\varphi \right), \tag{4.17}$$

where (ρ, φ) are now the new polar coordinates (see fig. 4.7), with $\rho = f\sin\vartheta$ and $\varphi = \phi$ and the old coordinate vectors $\hat{\boldsymbol{e}}_\vartheta$, $\hat{\boldsymbol{e}}_\phi$ understood as functions of the new variables. We have again neglected any polarisation effects at the interface, which in practice assumes that Fresnel losses are minimised by good anti-reflection coatings.

Applying the same argument in reverse, the field right after the second reference surface is then

$$\boldsymbol{\epsilon}_C(\vartheta', \phi') = -\sqrt{\frac{3}{8\pi}} \frac{\mathrm{e}^{ikf}}{f} \sqrt{\frac{\cos\vartheta'}{\cos\vartheta}} \left((\hat{\boldsymbol{e}}_\vartheta \cdot \hat{\boldsymbol{d}}) \hat{\boldsymbol{e}}_{\vartheta'} + (\hat{\boldsymbol{e}}_\phi \cdot \hat{\boldsymbol{d}}) \hat{\boldsymbol{e}}_{\phi'} \right), \tag{4.18}$$

now given in spherical coordinates (r', ϑ', ϕ'), with $\phi' = \varphi = \phi$ and $\frac{\sin\vartheta'}{\sin\vartheta} = \frac{f}{f'}$. In stating this, we have assumed the system to be aplanatic, that is, free from extra terms due to aberrations, and have omitted any common phase acquired due to the spatial (z) extent of the imaging system. It should be stressed that the two-step process $\boldsymbol{\epsilon}_A \mapsto \boldsymbol{\epsilon}_B \mapsto \boldsymbol{\epsilon}_C$ is

merely a model to motivate the transformation resulting from an ideal aplanatic optical system. In practice, the optical system will likely consist of a larger number of surfaces, the total effect of which is to approximate this ideal mapping between reference spheres.

Let $r' = (x', y', z')^T$ be a point in the image space. According to Wolf's formalism [Wol59], we can express the electric field amplitude ϵ' in that point as

$$\epsilon_\mathrm{D}(x', y', z') = \frac{\mathrm{i}k \mathrm{e}^{-\mathrm{i}kr'}}{2\pi} \iint\limits_{k_x{}^2 + k_y{}^2 \leq k^2} \frac{\epsilon_\mathrm{C}(k_x, k_y)}{k_z} \mathrm{e}^{\mathrm{i}(k_x x' + k_y y' + k_z z')} \, \mathrm{d}k_x \, \mathrm{d}k_y \qquad (4.19)$$

as long as r' is far from the reference surface (which is the case if r' is close to the focus, the case we are interested here).

To evaluate this, it is convenient to use cylindrical coordinates $r' = (\rho' \cos \varphi', \rho' \sin \varphi', z')^T$ in the image space. After substituting the field from eq. (4.18), the expression becomes

$$\epsilon_\mathrm{D}(\rho', \varphi', z') = -\sqrt{\frac{3}{8\pi}} \frac{\mathrm{i}kf'}{2\pi f} \mathrm{e}^{\mathrm{i}k(f-f')} \int\limits_0^{\vartheta_0'} \mathrm{e}^{\mathrm{i}kz' \cos \vartheta'} \sin \vartheta' \int\limits_0^{2\pi} \sqrt{\frac{\cos \vartheta'}{\cos \vartheta}} \left((\hat{e}_\vartheta \cdot \hat{d}) \hat{e}_{\vartheta'} + (\hat{e}_\phi \cdot \hat{d}) \hat{e}_{\phi'} \right)$$
$$\mathrm{e}^{\mathrm{i}k\rho' \sin \vartheta' \cos(\phi' - \varphi')} \mathrm{d}\phi' \mathrm{d}\vartheta', \qquad (4.20)$$

where the integration bound ϑ_0' corresponds to the circular aperture of the system. This is the final expression for the field at the given image space coordinates (ρ', φ', z'), which ultimately needs to be evaluated numerically.

To explicitly give the field in Cartesian coordinates, we expand the coordinate vectors as in eq. (4.13), and analogously for the primed coordinates. After substituting $\phi' = \phi$, $\sin \vartheta' = \sin \vartheta / m$ (and similarly $\cos \vartheta' = \sqrt{1 - (\sin \vartheta / m)^2}$), where we have defined the transversal magnification $m := f' / f$, and expressing the integration over ϕ using Bessel functions of the first kind, we are left with an integral over $\vartheta \in [0, \vartheta_0]$, where ϑ_0 is the boundary given by the object-side numerical aperture $\sin \vartheta_0$.

Finally, note that the constant phase factor has no clear interpretation here, as we have arbitrarily assumed that the common phase through the imaging system is zero. We thus drop it, that is, define

$$\epsilon' := -\mathrm{i} \, \mathrm{e}^{\mathrm{i}k(f'-f)} \epsilon_\mathrm{D}, \qquad (4.21)$$

and are left with the field of a dipole $\hat{d}$ in a point $r' = (\rho' \cos \varphi', \, \rho' \sin \varphi', \, z')^T$ near the focus

$$\epsilon'(r') = \sqrt{\frac{3}{8\pi}} \frac{k}{2m} \underbrace{\begin{pmatrix} I_0 + I_2 \cos(2\varphi') & I_2 \sin(2\varphi') & -\mathrm{i}I_1 \cos \varphi' \\ I_2 \sin(2\varphi') & I_0 - I_2 \cos(2\varphi') & -\mathrm{i}I_1 \sin \varphi' \\ -\mathrm{i}I_{z1} \cos \varphi' & -\mathrm{i}I_{z1} \sin \varphi' & -I_{z0} \end{pmatrix}}_{=: P} \hat{d}, \qquad (4.22)$$

where the azimuthal angle integrals (real-valued and positive for $z' = 0$) have been abbreviated as

$$I_0(\rho', z') := I\left[\left(1 + \cos\vartheta\cos\vartheta'\right) J_0\left(\frac{k\rho'}{m}\sin\vartheta\right)\right], \tag{4.23a}$$

$$I_1(\rho', z') := 2\, I\left[\sin\vartheta\cos\vartheta'\, J_1\left(\frac{k\rho'}{m}\sin\vartheta\right)\right], \tag{4.23b}$$

$$I_2(\rho', z') := I\left[\left(1 - \cos\vartheta\cos\vartheta'\right) J_2\left(\frac{k\rho'}{m}\sin\vartheta\right)\right], \tag{4.23c}$$

$$I_{z0}(\rho', z') := \frac{2}{m^2}\, I\left[(\sin\vartheta)^2 J_0\left(\frac{k\rho'}{m}\sin\vartheta\right)\right], \tag{4.23d}$$

$$I_{z1}(\rho', z') := \frac{2}{m^2}\, I\left[\sin\vartheta\cos\vartheta\, J_1\left(\frac{k\rho'}{m}\sin\vartheta\right)\right], \text{ with} \tag{4.23e}$$

$$I[g] := \int_0^{\vartheta_0} \sqrt{\frac{\cos\vartheta}{\cos\vartheta'}}\, g(\vartheta)\, e^{ikz'\cos\vartheta'}\sin\vartheta\, d\vartheta, \quad \cos\vartheta' := \sqrt{1 - \frac{(\sin\vartheta)^2}{m^2}}. \tag{4.23f}$$

For later notational convenience, we have defined P as the "propagator" matrix $P(\rho', \varphi', z')$ that describes the transformation from dipole $\hat{\boldsymbol{d}}$ to image field $\boldsymbol{\epsilon}'(\boldsymbol{r}')$.

Note that the image-space field does not directly depend on f' and f, only on their ratio, the magnification m. This is expected, as we assume that the focal lengths are much larger than the wavelength, and as such the boundary conditions on the aperture stop can be neglected.

Also note that the z component of eq. (4.22), that is, the longitudinal field, does not quite vanish. This is to be expected, as $\boldsymbol{\epsilon}'(\boldsymbol{r}')$ is a rigorous solution of the source-free Maxwell equations, i.e. $\nabla \cdot \boldsymbol{\epsilon}' = 0$, so only plane waves are strictly transversal. The longitudinal field amplitudes scale like the inverse square of the magnification, though, so for large m, the longitudinal contribution can be safely neglected.

4.2.3   Atomic dipole images

We can now calculate the electric fields generated by the atomic dipole in the focal plane of the imaging system by inserting the appropriate dipole moments into eq. (4.22). For all the figures shown in this section, we will use dipole moments rescaled according to the Clebsch–Gordan coefficients,

$$\tilde{\boldsymbol{d}}_{\sigma^+} = (1, 0, -i)^T/\sqrt{3}, \quad \tilde{\boldsymbol{d}}_{\pi} = (0, 1, 0)^T/\sqrt{3}, \tag{4.24}$$

for ease of visual interpretation. For maximal entanglement, the two fields should then have the same absolute value. The fractional collection efficiency, that is, the probability of collecting a photon after excitation, corresponds directly to the (integral of) the sum of
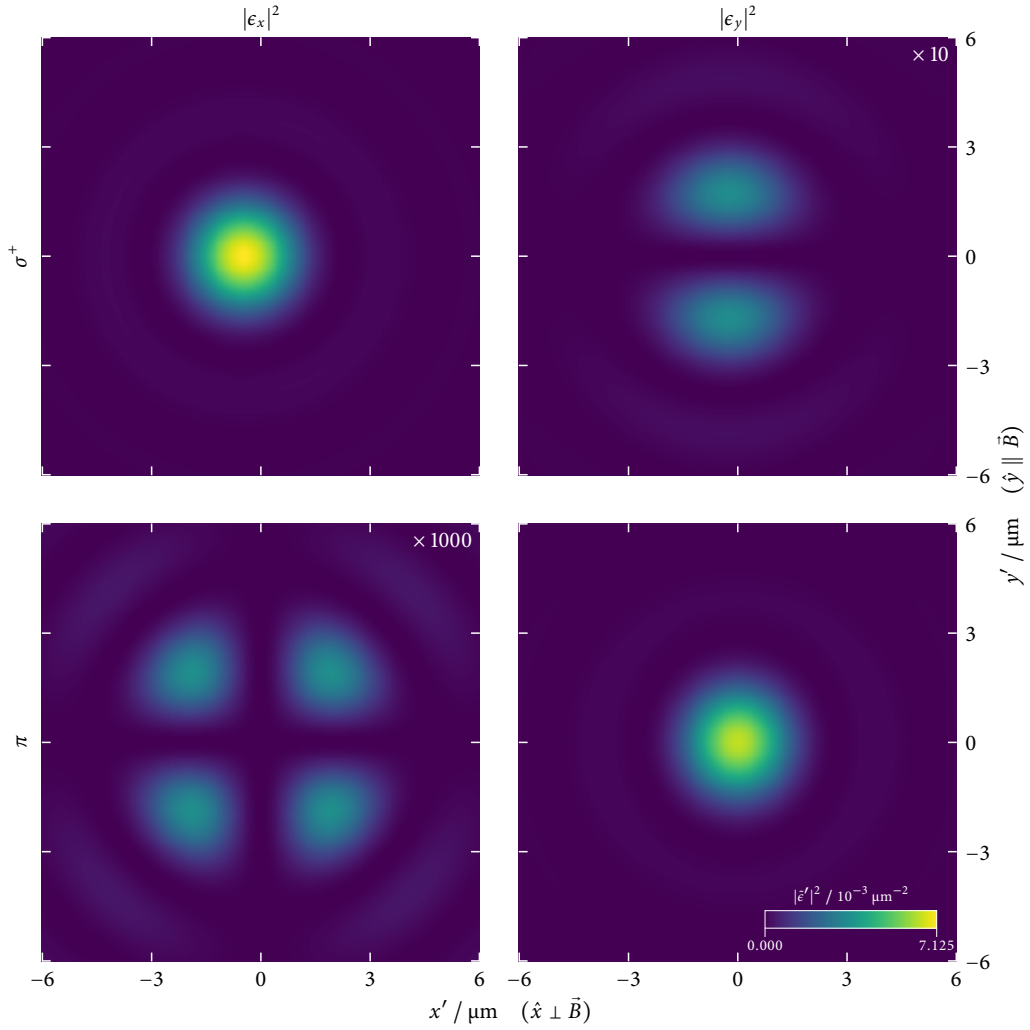
Figure 4.8: Electric field intensity distribution for π and $\sqrt{2}\,\sigma^+$ dipoles in the image plane of a numerical aperture $\sin\vartheta_0 = 0.6$ objective with magnification $6.7\times$ ($f'/f = 1/0.15$). The same colour scale is used for all plots, with the suppressed polarisations multiplied by the given factors for visual clarity.

their squared absolute values. Again in terms of the quantities from eq. (4.23), the fields then read

$$\tilde{\boldsymbol{\epsilon}}'_{\sigma^+} = \sqrt{\frac{1}{8\pi}}\frac{k}{2m}\begin{pmatrix} I_0 - I_1\cos\varphi' + I_2\cos(2\varphi') \\ -I_1\sin\varphi' + I_2\sin(2\varphi') \\ \mathrm{i}\,I_{z0} - \mathrm{i}\,I_{z1}\,\cos\varphi' \end{pmatrix}, \quad \tilde{\boldsymbol{\epsilon}}'_{\pi} = \sqrt{\frac{1}{8\pi}}\frac{k}{2m}\begin{pmatrix} I_2\sin(2\varphi') \\ I_0 - I_2\cos(2\varphi') \\ -\mathrm{i}\,I_{z1}\sin\varphi' \end{pmatrix}.$$

$$(4.25)$$

In the limit of small collection apertures, that is, $\vartheta_0 \ll \pi/2$, the terms integrating over Bessel functions J_1 and J_2 are negligible, and I_{z0} is also suppressed due to the $(\sin\vartheta)^2$ factor. For the σ^+ transition, only the $\hat{\boldsymbol{x}}$ portion of the dipole moment perpendicular to the collection axis contributes, and we have recovered the on-axis case, where σ^+ and π give rise to orthogonal photon polarisations of equal amplitude.

For large numerical apertures on the collection side, this is only still true exactly on
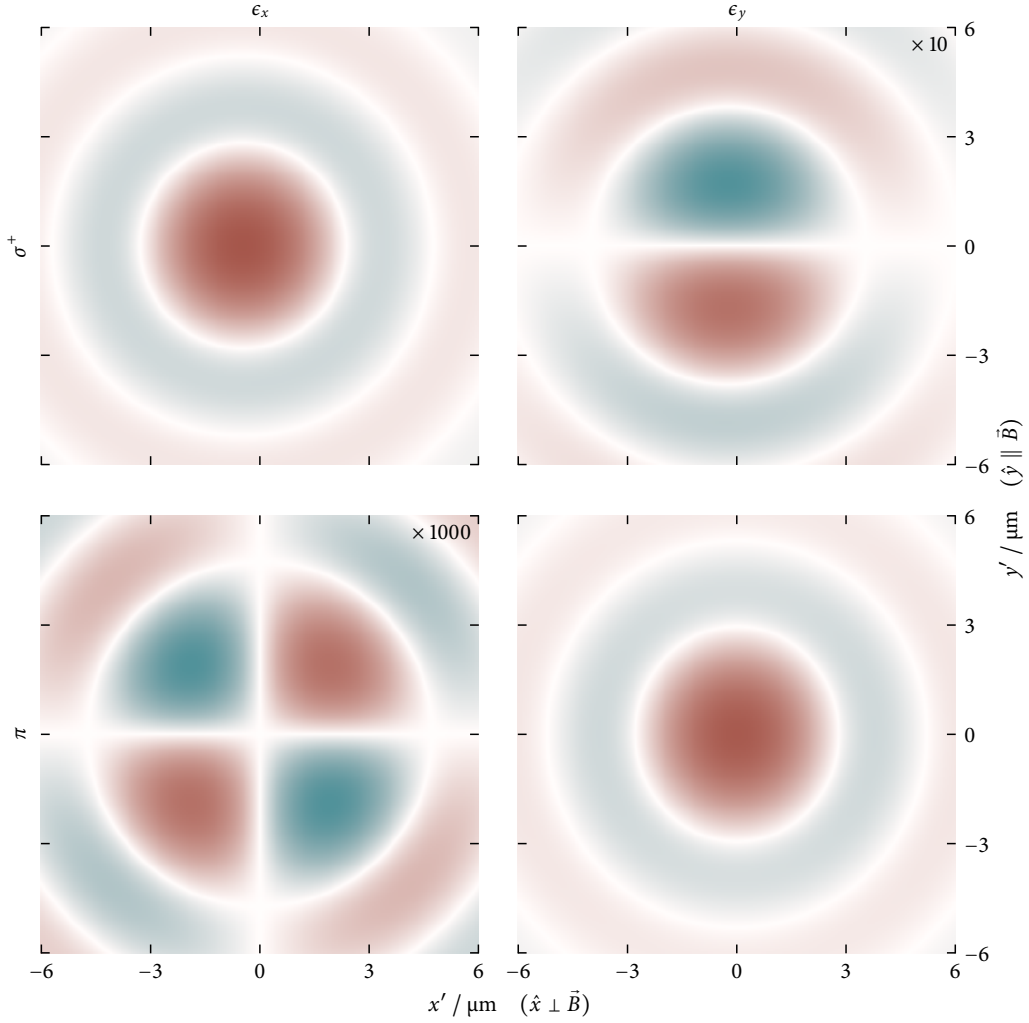
Figure 4.9: Electric field distribution for π and $\sqrt{2}\,\sigma^+$ dipoles in the image plane of a numerical aperture $\sin\vartheta_0 = 0.6$ objective with magnification $6.7\times$ ($f'/f = 1/0.15$). Colours represent the phase of the field, which (in the chosen convention) is real-valued here in the focal plane, brightness the (inverse) amplitude. The same colour scale is used for all plots, with the suppressed polarisations multiplied by the given factors for visual clarity.

the optical axis; the two polarisations are indeed no longer orthogonal off it. The case for numerical aperture $\sin\vartheta_0 = 0.6$ and magnification $m = 1/0.15$ is shown in figs. 4.8 and 4.9.

The longitudinal field components (ϵ_z') still contribute less than 5×10^{-5} to the total intensity (a corresponding figure is thus omitted); this would also be true for even larger collection angles if the image-side numerical apertures is kept at the fairly limited value for mode-matching typical single-mode fibres (here $\sin\vartheta_0/m = 0.09$) – the paraxial approximation is still fairly good on the image side.

The other transversal terms are now no longer negligible, though, in particular the I_1 term from the $\hat{z}$ dipole moment component of the σ^+ transition parallel to the collection axis. Note, however, that all these contributions are periodic in φ' (i.e. symmetric about
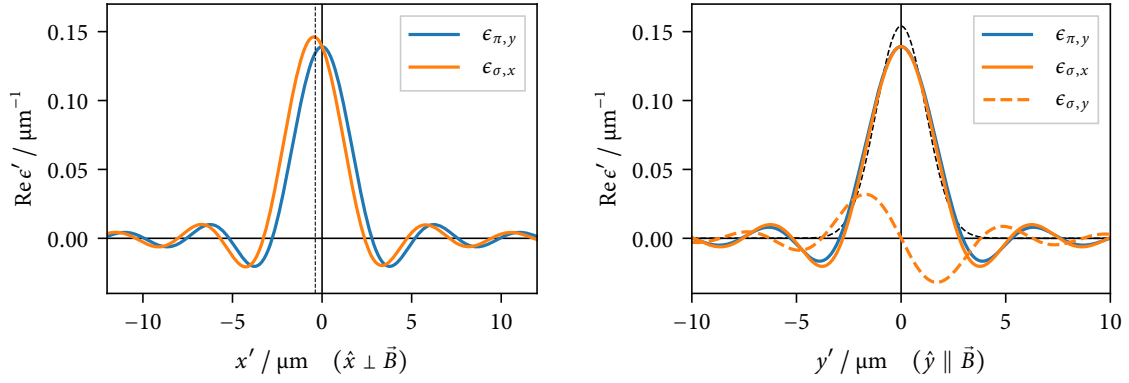
Figure 4.10: Non-zero electric field components along the coordinate axes through the focus of $\hat{\boldsymbol{\pi}}$ and $\sqrt{2}\hat{\boldsymbol{\sigma}}^+$ dipoles, in the image plane of a numerical aperture $\sin\vartheta_0 = 0.6$ objective with magnification $6.7\times$ ($f'/f = 1/0.15$). *Left:* Slice perpendicular to the magnetic field direction. Here, both polarisations are orthogonal, but the $\hat{\boldsymbol{\sigma}}^+$ emission maximum is shifted due to photon orbital angular momentum effects. The dashed line shows the $\delta x = 379$ nm low-NA prediction for the σ^+ image centroid from ref. [Ara+18]. *Right:* Slice parallel to the magnetic field direction. Off the $x' = 0$ plane, the fields are no longer orthogonal, but the ϵ_y' component generated by σ emission is odd in x'. The black dashed line shows a Gaussian mode with $\frac{1}{e^2}$-radius 1.67 μm for comparison (see §4.2.4).

the optical axis) in such a way that the average across circles $\varphi' \in [0, 2\pi)$ is zero. This is particularly evident in fig. 4.9; we shall make use of this in the next section.

Shift in apparent dipole position for $\sigma^\pm$ emission. Figures 4.8 and 4.9 reveal another interesting feature of the dipole emission geometry: The σ^+ peak appears to be shifted with respect to the coordinate origin in the optical axis! This is readily visible in fig. 4.10, which shows one-dimensional slices through the field amplitudes along the coordinate axes (with the convention chosen here, the transversal fields are real in the image plane, cf. eq. (4.21)). This phenomenon is not related to defocus; the shift appears even in the focal plane $z' = 0$, such as for the fields shown here.

Mathematically, this can be seen as a consequence of the longitudinally-oriented $\hat{\boldsymbol{z}}$ component of the σ^+ dipole to ϵ_x', the $I_1 \cos\varphi'$-contribution of which overlaps destructively with the I_0 term from the $\hat{\boldsymbol{x}}$ dipole component for $x' > 0$, and constructively for $x' < 0$ (with opposite signs for a σ^- transition).

Physically, this can be interpreted as a consequence of the conservation of angular momentum: In the plane perpendicular to the magnetic field, σ^+ emission is linearly polarised. As such, the angular momentum is carried away not in the form of photon helicity, but must rather be encoded as orbital angular momentum of the photon mode geometry. The resulting spiralling wavefronts give rise to a phase ramp across the aperture of the collection objective, or a positional shift in the image plane. This has been discussed for the case of a classical dipole scatterer in ref. [SD06], and more recently in ref. [Ara+18]. In the latter work, the quantum case is examined explicitly, where imaging through an aperture can be interpreted as a weak measurement of the photon momentum. In the limit of low

numerical aperture, the shift in image centroid expected for a $\sigma^\pm$ transition is shown to be

$$\delta x \approx \mp \frac{m}{k} \frac{1}{1 + \frac{(\sin \vartheta_0)^2}{2}}, \qquad (4.26)$$

which is only slightly exceeded for the imaging parameters considered here, as shown in fig. 4.10.

Larger collection angles. While only the NA 0.6 case was shown in the diagrams here, all the above symmetry considerations still apply even in the limit of large numerical apertures, i.e. as the collection cone half-angle approaches $\pi/2$. The relative strength of the I_2 second-order Bessel function integral terms, which only contribute on the $\approx 2 \times 10^{-3}$ level for NA 0.6, will increase, though. In particular, this will cause appreciable polarisation perpendicular to the magnetic field to be observed also for π transitions, and the $\hat{y}$-polarised image for π transitions to take on an elongated shape. The latter effect is the dual of what one observes when strongly focussing a linearly polarised laser beam, which, in the plane wave limit as well as considering a Gaussian profile, has been known for considerable time [RW59; NH06].

4.2.4 *Single-mode fibre coupling*

We now turn our attention to the case where the electric fields are not directly observed, but an optical fibre is placed in the image plane instead. For an appropriate choice of parameters, a cylindrically symmetric optical fibre with a single step in refractive index along its radial profile supports only a single propagating mode of either polarisation. Assuming that it is long enough for evanescently coupled light to have decayed away, it thus acts as a spatial mode filter, with the allowed mode being well-approximated by a transversal Gaussian profile. For the sake of convenience, we choose a normalisation similar to that we chose for the atomic modes, $\iint_{\mathbb{R}^2} |\hat{\epsilon}_F(x, y)|^2 \mathrm{d}x \, \mathrm{d}y = 1$. The modes of a fibre aligned with the $\hat{z}$ axis are then given by

$$\epsilon_{F,k}(\rho', \varphi') = \underbrace{\sqrt{\frac{2}{\pi s^2}} \, e^{-(\rho'/s)^2}}_{=:u(\rho')} \hat{k}, \qquad (4.27)$$

where the parameter s describes the mode-field radius (giving the $\frac{1}{e^2}$-intensity radius of the Gaussian beam), and where $\hat{k}$ enumerates some arbitrary orthonormal basis of the plane perpendicular to the fibre axis. The coordinate axes $\hat{x}$ and $\hat{y}$ are an expedient choice for the polarisation modes here, and we note that the following considerations would be equally valid for different radial mode profiles $u(\rho')$.

In the quantum setting, this can be viewed as a projection of the spatial degrees of freedom of the photon field onto a two-dimensional system describing the photon polarisation. The joint (unnormalised) ion–photon state at the output of the single-mode fibre is thus simply

$$|\psi\rangle_{IP} = \tilde{\epsilon}_{\sigma,x} |0\rangle_I |H\rangle_P + \tilde{\epsilon}_{\sigma,y} |0\rangle_I |V\rangle_P + \tilde{\epsilon}_{\pi,x} |1\rangle_I |H\rangle_P + \tilde{\epsilon}_{\pi,y} |1\rangle_I |V\rangle_P, \qquad (4.28)$$

where the amplitudes are given by the usual inner product of dipole and fibre spatial modes as complex functions on $\mathbb{R}^2$, i.e. the overlap integrals

$$\tilde{\epsilon}_{d,k} = \iint\limits_{\mathbb{R}^2} \overline{\epsilon_{F,k}(r)} \cdot \epsilon'_{\hat{d}}(r) \, d^2r. \qquad (4.29)$$

In general, these integrals will have to be evaluated numerically. However, for the case where the fibre is aligned with the optical axis, i.e. the coordinate systems in eqs. (4.22) and (4.27) coincide, the structure greatly simplifies due to symmetry: Writing

$$\tilde{\epsilon}_{d,k} = \int\limits_0^\infty \int\limits_0^{2\pi} \overline{\epsilon_{F,k}(\rho',\varphi')} \cdot \epsilon'_{\hat{d}}(\rho',\varphi',z') \, d\varphi' \, \rho' \, d\rho' = \int\limits_0^\infty u(\rho') \int\limits_0^{2\pi} (\epsilon'_{\hat{d}})_k(\rho',\varphi',z') \, d\varphi' \, \rho' \, d\rho', \qquad (4.30)$$

all the non-I_0 terms in eq. (4.22) cancel on evaluating the φ' integral as the fibre mode profile u has no angular dependence, and we can write the resulting amplitudes more compactly, introducing the centrally symmetric part of the propagator $\tilde{P}$, as

$$\tilde{\epsilon}_d = \sqrt{\frac{3}{8\pi}}\frac{k}{2\,m} \underbrace{\begin{pmatrix} \int_0^\infty u(\rho') \, I_0(\rho',z') \, \rho' \, d\rho' & 0 & 0 \\ 0 & \int_0^\infty u(\rho') \, I_0(\rho',z') \, \rho' \, d\rho' & 0 \end{pmatrix}}_{=:\tilde{P}} \hat{d}. \qquad (4.31)$$

If $u(\rho')$ is the profile of a Gaussian mode as defined in eq. (4.27), we can also analytically evaluate the integral over ρ', giving

$$\tilde{P} = \begin{pmatrix} \sqrt{3}\tilde{I}_0 & 0 & 0 \\ 0 & \sqrt{3}\tilde{I}_0 & 0 \end{pmatrix} \qquad (4.32)$$

with

$$\tilde{I}_0 := \frac{s\,k}{2m} \int\limits_0^{\vartheta_0} \sqrt{\frac{\cos\vartheta}{\cos\vartheta'}} \, (1 + \cos\vartheta\cos\vartheta') \, e^{-(\frac{s\,k}{2m})^2(\sin\vartheta)^2} \, e^{ikz'\cos\vartheta'} \sin\vartheta \, d\vartheta, \qquad (4.33)$$

where $\cos\vartheta' = \sqrt{1 - (\sin\vartheta/m)^2}$ as before. The (unnormalised) ion–photon state from eq. (4.28) is thus simply

$$|\psi\rangle_{IP} = \tilde{I}_0 \left(|0\rangle_I |H\rangle_P + |1\rangle_I |V\rangle_P\right). \qquad (4.34)$$

After post-selection on the presence of a photon, this is a maximally entangled state! This is a consequence of the radial symmetry, and as such valid *independently* of the numerical
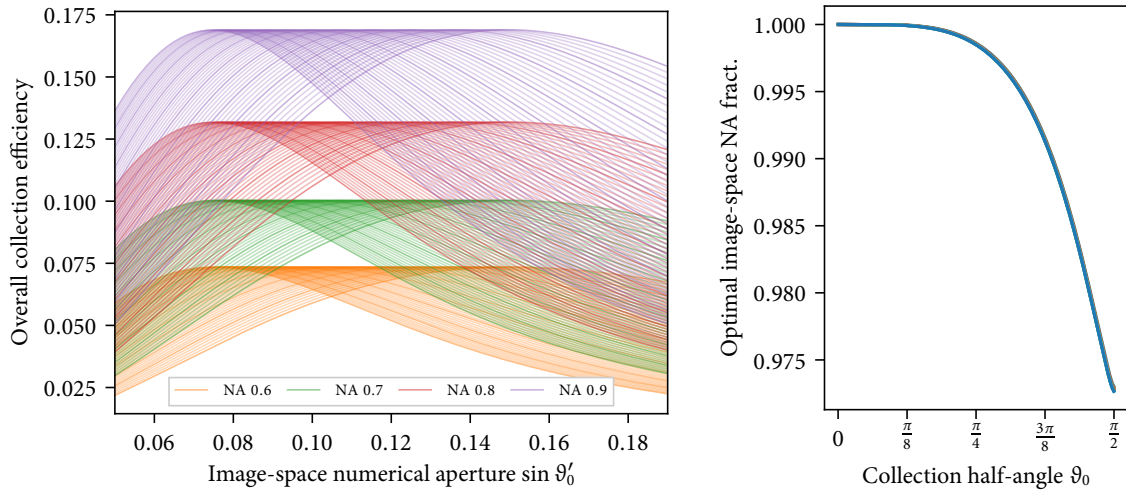
Figure 4.11: Collection efficiencies into single-mode fibre. *Left:* Overall collection efficiency from atom into single-mode fibre at 422 nm, as a function of the fibre-side numerical aperture. Colours represent four different ion-side collection angles as per the legend. In each case, the coupling efficiencies are shown for 51 different fibre mode field radii from 1 μm to 2 μm, with heavy lines marking steps of 0.1 μm. For each mode field radius, the maximum in collection efficiency is attained at an approximately constant image-space numerical aperture, regardless of the collection-side angle. *Right:* For mode-field radii across the same range, the image-space numerical apertures at which maximum fibre-coupling is obtained (given as a fraction of the respective value in the limit of zero ion emission collection angle) is nearly constant.

aperture, or the radial mode profile of the fibre; as these parameters are varied, only the common $\tilde{I}_0$ prefactor changes. Note again that we have assumed that the fibre mode is centred on the collection axis; the consequences arising from misalignment are discussed in the next section.

To maximise the collection efficiency $\eta \coloneqq 2\,|\tilde{I}_0|^2$ for a given set of parameters, one would typically choose a certain numerical aperture $\sin\vartheta_0$ for the collection optics based on engineering constraints, and the mode radius s would be given by that of typical silica single-mode fibres at the target wavelength. The optimal magnification m of the imaging system would then be chosen numerically to maximise eq. (4.33).

To good approximation, maximal coupling will be attained at the point where the image-side numerical aperture $m\sin\vartheta_0$ best matches the divergence angle of the fibre mode. There will be some numerical constant relating this value to the "core NA" often catalogued by fibre manufacturers due to the difference in mode profiles. This holds across different fibre mode-field radii; the residual dependence on the collection angle ϑ_0 is small, the optimum reducing by 2.5% in the limit of collecting whole half-space, as shown in fig. 4.11.

For numerical apertures $\sin\vartheta_0 = 0.6$ and $\sin\vartheta_0' = 0.09$, the optimal mode-field radius at 422 nm is $s = 1.67$ μm. The corresponding mode amplitude was depicted in the earlier fig. 4.10. Clearly, the overlap is not perfect, owing to the difference between the Gaussian fibre profile and the Airy-like ion emission mode.
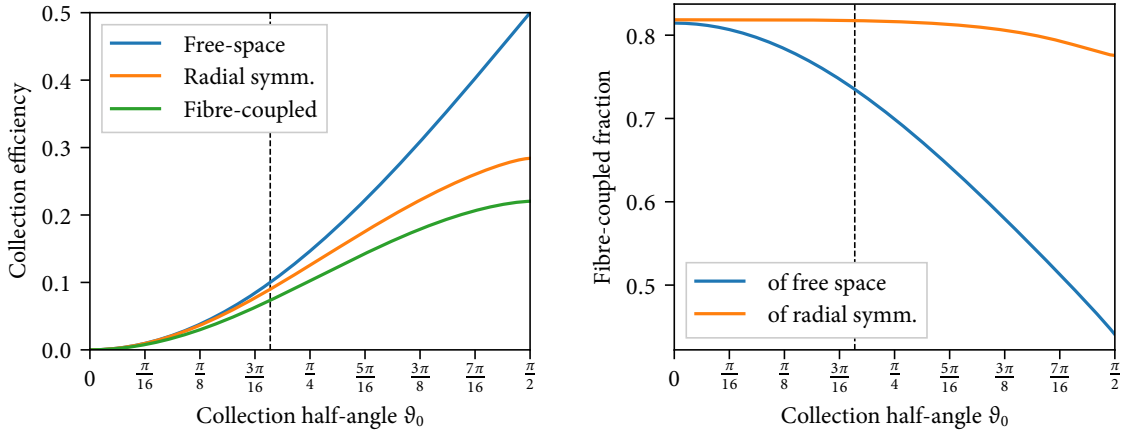
Figure 4.12: Achievable efficiencies when coupling ion radiation into a Gaussian single-mode fibre. Dashed lines show NA 0.6, where the free-space collection efficiency is 0.1, and the absolute fibre-coupled collection efficiency is 0.074. This is 74 % of the free-space fraction, and 81 % of the radially symmetric part.

The maximum achievable collection efficiencies for various collection angles are shown in fig. 4.12. The emission from $J = 1/2 \leftrightarrow J' = 1/2$ transitions is spherically symmetric, so the fraction of ion emission collected in free space is simply the solid angle $\eta_{\mathrm{fs}} := \sin(\vartheta_0/2)^2$ (which of course can easily be verified by integrating over the expression from eq. (4.3)). As only the I_0 terms are coupled into the fibre, the relative fibre-coupling efficiency η/η_{fs} falls towards wider collection angles, as shown in fig. 4.12.

By itself, this is neither surprising, nor undesired: as the numerical aperture increases and the collection cone includes radiation from closer to the magnetic field axis, $\sigma^{\pm}$ radiation will gain in (relative) strength. If all this radiation was coupled into the fibre, the balance between σ^+ and π would be upset, shifting the resulting state away from being maximally entangled. The more salient comparison to quantify the mode-matching quality is thus against the "useful" components, given by the I_0 terms in eq. (4.22), which are of equal strength for σ^+ and π. Since this excludes any contribution from the $\hat{z}$-oriented component of the σ^+ emission, the fraction $\eta_{I_0}/\eta_{\mathrm{fs}}$ similarly falls off heavily towards larger collection angles.

The remaining loss in coupling due to mismatch of the Gaussian fibre mode with the symmetric part of the ion emission can then be quantified by η/η_{I_0}, also shown on the right side of fig. 4.12. For moderate collection angles, the image-plane fields do not deviate much from the low-NA limit (and from what would be the point-spread function of the imaging system), and the fractional coupling efficiency remains almost constantly at the 81.8 % for low numerical apertures, only falling below 80 % at NA 0.96, and reaching 77.6 % as ϑ_0 approaches $\pi/2$. For realistic collection angles, an ≈ 19 % improvement in collection efficiency could thus be realised by tailoring the shape of the fibre mode to match the perpendicular dipole image.

4.3   Ion–photon entanglement imperfections

To arrive at the simple form from eq. (4.31) for the fibre-coupled emission amplitudes, a number of assumptions were made regarding the system geometry: Ion and fibre tip being exactly in the respective foci of the imaging system, the optical axis being exactly perpendicular to the magnetic field and the fibre being aligned along the same, the ion being a stationary dipole emitter, etc. While the theoretical prediction of unit ion–photon entanglement fidelity is encouraging, it will no longer hold if any of these conditions are violated. For the design and operation of an experiment making use of this scheme, it is crucial to have an estimate for the sensitivity of expected fidelity and collection efficiency to the various error sources.

In this section, we will quantitatively consider the effect of a number of imperfections. While the ion–photon state can still easily be calculated numerically from the full expressions for field amplitude and fibre overlap integrals, it is not straightforward to obtain analytical results outside the low-NA limit for most calculations. Throughout this section, we will thus consider the particular case of 422 nm radiation imaged by a NA 0.6 objective with magnification $f'/f = 1/0.15$, and the corresponding optimal Gaussian fibre mode of $\frac{1}{e^2}$-radius $s = 1.67\,\mu\text{m}$. These are the parameters most relevant to the experiments discussed in the rest of this thesis, but should also be fairly typical for other atom-photon interface experiments of this kind.

To provide a simple benchmark regarding the relative sensitivity towards the discussed effects, I will also state the amount of imperfection at which the fidelity (fully entangled fraction) of the generated state has dropped by an additional error $\mathcal{E} = 0.5\,\%$ over the theoretical perfect Bell state, and also where the collection efficiency has dropped by 5 %. These thresholds are chosen fairly arbitrarily, but any experiment aiming to reach state-of-the-art atom-light entanglement fidelities should be constructed to restrict most error sources to comfortably below the given figures.

4.3.1   Positional misalignment

Given that the most straightforward way to focus ion emission into a fibre in the experiment is to use a macroscopic lens objective outside the vacuum system holding the ion trap, which will thus not be directly registered to the ion position mechanically and involve focal lengths on the order of many centimetres, the sensitivity to small positional misalignments is experimentally very relevant.

We can readily model a change Δr in the fibre position by modifying the coordinate origin of the fibre mode, that is, making the replacement $\epsilon_{F,k}(r) \mapsto \epsilon_{F,k}(r - \Delta r)$ in eq. (4.29). This is most naturally discussed in two distinct cases, that of transversal displacements (i.e. perpendicular to the optical axis), and that of longitudinal displacements along the optical axis (i.e. defocus).
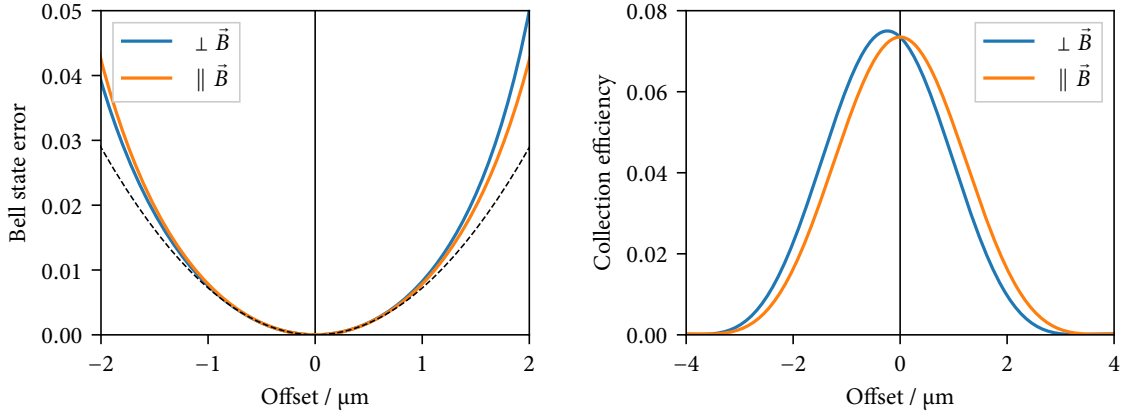
Figure 4.13: Fibre-coupled performance under transversal misalignment of the collection fibre. *Left:* For small offsets, both directions affect the fully entangled fraction of the state almost equally; the dashed line shows the numerically obtained approximation $7.23 \times 10^{-3}\,\mu\mathrm{m}^{-2}\,\delta r^2$. *Right:* The sensitivity of the overall fraction of collected photons is similar in both directions, but perpendicular to the magnetic field, the maximum is offset due to the orbital angular momentum shift in the σ^+ image pattern.

The effect of shifts in position on the object, that is, ion trap side can then be inferred using the geometrical imaging relations of the optical system, as discussed below.

Transversal misalignment. If the fibre tip is moved off the optical axis, the coupling of the I_1 and I_2 terms in eq. (4.25) to the fibre mode no longer vanishes. As $\|I_1\| \gg \|I_2\|$, the ion–photon state fidelity will predominantly be affected by the former, corresponding to the contribution of the $\hat{z}$-oriented part of the σ^+ dipole. For shifts along $\hat{x}$, that is, perpendicular to the magnetic field, the extra contribution from the $-I_1 \cos\varphi'$ term in $\epsilon'_{\sigma^+,x}$ will upset the balance between the π and σ^+ components of the entangled state, whereas for shifts along the $\hat{y}$ direction of the magnetic field, the $-I_1 \sin\varphi'$ term in $\epsilon'_{\sigma^+,y}$ will lead to a non-zero amplitude for detecting a $|V\rangle$ photon from a σ decay.

The result of the numerical calculations for displacements along the coordinate axes is shown in fig. 4.13. For small shifts, the resulting loss in fidelity is almost equal for both directions, and for the given parameters can be approximated as

$$\mathcal{F}_{\text{shifted}}(\delta r) \approx 1 - 7.23 \times 10^{-3}\,\mu\mathrm{m}^{-2}\,\delta r^2.$$

The fidelity thus reaches 99.5 % at 832 nm image-side displacement, where the collection efficiency (in the parallel case) will have fallen to 78 % of the peak value. To not lose more than 5 % in collection efficiency, a shift of at most ≈ 380 nm can be tolerated.

Note that the total collection efficiency actually increases for small shifts towards the orbital-angular-momentum-displaced σ^+ emission peak (see last section). When the fibre position is adjusted in the experiment to maximise the amount of fluorescence collected from the ion, it is thus preferable to either consider only the $|V\rangle$ photons corresponding to the π transition where the coupling maximum actually coincides with the fibre axis, or
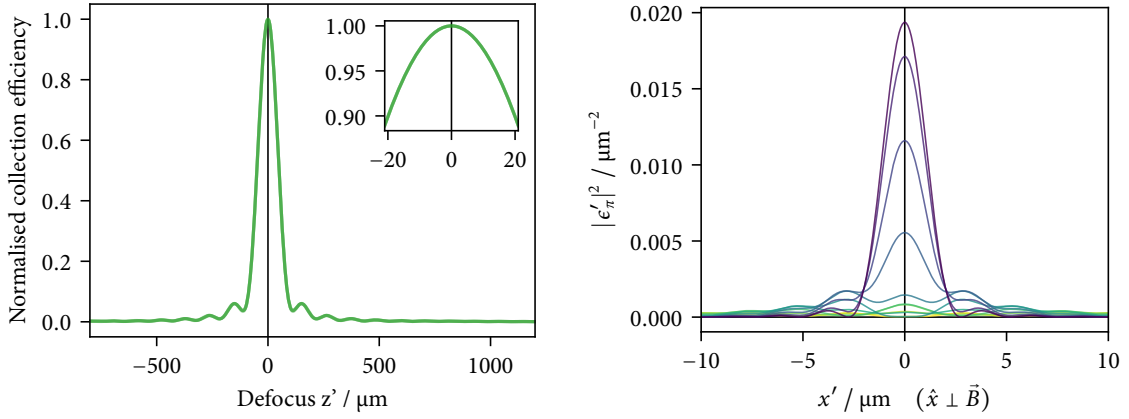
Figure 4.14: Fibre-coupling efficiencies under defocus (fidelity is not affected). *Left:* Coupling efficiencies as a function of the image-space displacement from the focus. As the object-side numerical apertures varies between 0.1 and 0.9 (while keeping the image-side NA constant), the predicted coupling efficiencies agree to within the less than width of the displayed line. Inset shows a close-up of the maximum on the same axes. *Right:* A series of transversal slices along the $\hat{x}$ axis of the image field. Curves shown are for z' values from 0 to 200 μm, in steps of 20 μm, as marked by the colour gradient.

to use steady-state cooling fluorescence where σ^+ and σ^- are emitted in equal amounts (although the loss in fidelity from this would be fairly small, at 4×10^{-4}).

In the treatment covered here, we have consistently assumed the ion to be located exactly at the centre of the object-side reference sphere. For small displacements, however, we can make use of the geometric relations of the imaging system – the fact that m gives the transversal magnification – to obtain the image-side field after a small displacement of the dipole emitter to a point $(x, y, 0)^T$ a small distance from the coordinate origin as $\epsilon'(x' + m\,x, y' + m\,y, z')$. This is of course only true if the imaging system used is also well-corrected for points off the axis, and the shift does not yet lead to significant asymmetric clipping of the effective collection cone.[9] The $F_{\text{shifted}} = 0.995$ and $\eta_{rel} = 95\,\%$ thresholds for shifts on the ion side are then $\delta r = 125$ nm and $\delta r = 57$ nm, respectively.

Defocus. The behaviour under translation along the $\hat{z}$ direction is qualitatively different to the transversal directions, as it keeps the angular symmetries of the system intact. As such, the fidelity of the resulting ion–photon state is unaffected. The overall collection efficiency suffers, however, as the mode optimised for matching a small single-mode fibre will be quite divergent. The behaviour for the above model system parameters is shown in fig. 4.14; to be within 5 % of optimal coupling, a shift of at most $\delta z' = 14$ μm can be tolerated.

For a shift in the ion position from the origin of the reference sphere to $(0, 0, z)^T$, the image fields now shift according to the longitudinal magnification m^2, that is, are given by $\epsilon'(x', y', z' - m^2 z)$. The corresponding 95 % efficiency threshold for shifts in the ion

---

[9]The limits of this approximation heavily depend on the lens design; for the NA 0.6 objective used here, the largely unaberrated field of view is some 100 μm.
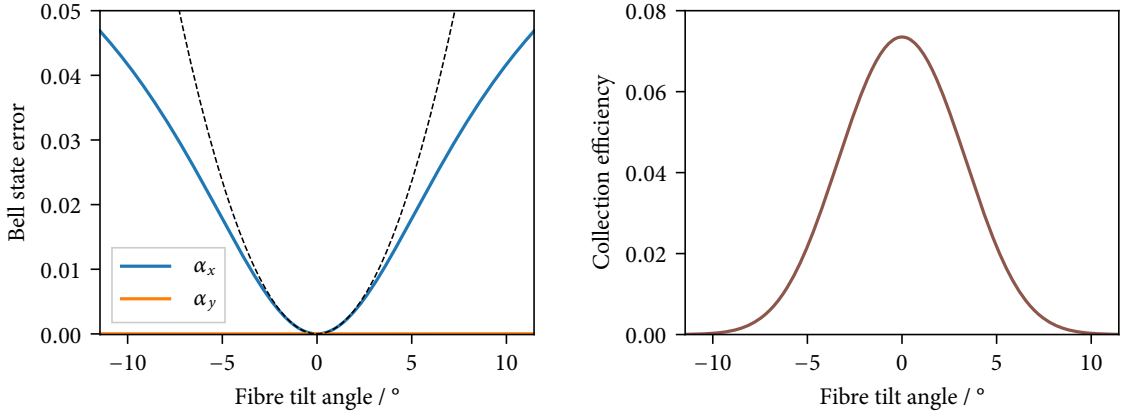
Figure 4.15: Ion-photon interface performance for angular misalignments of the single-mode fibre with the collection axis. *Left:* Tilts around the $\hat{y}$ axis parallel to the field do not introduce appreciable errors, while tilts around the $\hat{x}$ axis perpendicular to the field do. The dashed line shows a quadratic approximation for small angles $(3.1 \cdot (\alpha_x / \text{rad})^2)$. *Right:* Tilts in either direction equally reduce the overall coupling of photons into the single-mode fibre.

position is at $z = \pm 318\,\text{nm}$. While this is less stringent than the transversal position in absolute terms, the change occurs over the potentially significant front focal length f of the imaging system.

By way of example, $f = 20.9\,\text{mm}$ for the objective lenses used in the experimental section of this thesis, giving a fractional stability requirement of 15×10^{-6}. Typically, non-magnetic steel alloys are used in the construction of atom optics experiments, which tend to have quite a large coefficient of linear thermal expansion; $\alpha = 16 \times 10^{-6}\,\text{K}^{-1}$ for grade 316 austenitic stainless steel [CW86]. This implies that the temperature of a typical apparatus should be kept stable to better than $1\,\text{K}$, even if the parts that define the transversal alignment were matched in thermal expansion behaviour.

4.3.2   *Fibre tilt*

Angular misalignment of the collection fibre, that is, tilts away from the optical axis, can be modelled in much the same way as positional misalignment by appropriately modifying the coordinate definitions in eq. (4.29). Rotations of the fibre along its axis are inconsequential in terms of the fully entangled fraction, so we choose to represent arbitrary tilts as a succession of rotations $R_x(\alpha_x) \circ R_y(\alpha_y)$, where R_x and R_y can be expressed by the usual SO(3) rotation matrices

$$R_x(\varphi) := \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\varphi & -\sin\varphi \\ 0 & \sin\varphi & \cos\varphi \end{pmatrix}, \quad R_y(\varphi) := \begin{pmatrix} \cos\varphi & 0 & \sin\varphi \\ 0 & 1 & 0 \\ -\sin\varphi & 0 & \cos\varphi \end{pmatrix}. \tag{4.35}$$
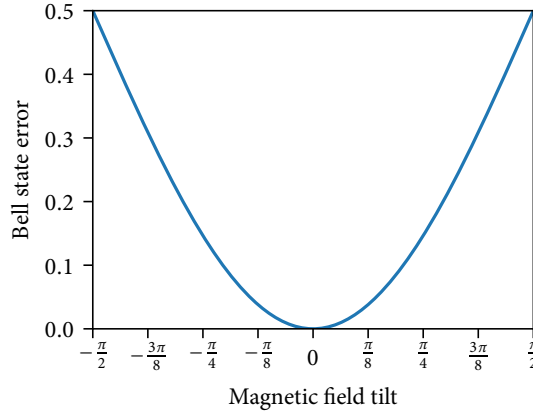
Figure 4.16: Loss in ion–photon fidelity introduced by tilts of the static magnetic field out of the plane perpendicular to the collection axis. The fraction of coupled photons remains unaffected.

Defining the rotated unit vectors corresponding to the direction of the fibre polarisation modes (which also span the plane defining the fibre face) as

$$\hat{e}_H = R_x(\alpha_x)\,R_y(\alpha_y)\,\hat{x}, \quad \hat{e}_V = R_x(\alpha_x)\,R_y(\alpha_y)\,\hat{y}, \tag{4.36}$$

the integration over the fibre face then reads

$$\tilde{\epsilon}_{d,j} = \sqrt{\frac{2}{\pi\,s^2}} \int\limits_{-\infty}^{\infty} \int\limits_{-\infty}^{\infty} e^{-(h^2+v^2)/s^2}\,\epsilon'_{\tilde{d}}(h\,\hat{e}_H + v\,\hat{e}_V)\cdot\hat{e}_j\,\mathrm{d}h\,\mathrm{d}v, \quad j \in \{H,V\}. \tag{4.37}$$

The resulting performance characteristics are shown in fig. 4.15. A strong asymmetry is visible in the behaviour: Tilting the fibre around the axis parallel to the field to very good approximation now just uniformly degrades coupling to the fibre due to what in the paraxial approximation is a phase ramp across the fibre phase. Before the loss in fidelity exceeds 1×10^{-5}, the total collection efficiency has already dropped to $< 0.1\,\%$ of the aligned case. Tilts around the axis perpendicular to the field decrease the overall collection efficiency in the same way. Now, however, the phase ramp breaks the symmetry which prevents the σ radiation to coupling to V polarisation, leading to a large polarisation mixing error. For small angles (given in radian),

$$\mathcal{F}_{\text{tilt}}(\alpha_x) = 1 - 3.1 \cdot \alpha_x{}^2 - \Omega(\alpha_x{}^4). \tag{4.38}$$

$F_{\text{tilt}} > 0.995$ is reached for tilt angles $\alpha_x < 2.3°$; the $\eta_{rel} = 0.95$ threshold is at $\alpha_j = 1.0°$ ($j = x, y$).

4.3.3 Magnetic field misalignment

Another case of angular misalignment is given when the direction of the magnetic field defining the atomic quantisation axis (through the Zeeman effect) differs from the intended orientation. If the magnetic field is rotated in the plane perpendicular to the collection axis, only the basis in which the state is generated changes; the state is still maximally

entangled. Rotations around the $\hat{x}$ axis, which leave the collection axis no longer orthogonal to the field, lead to an error, however.

Assuming the fibre is still well-aligned with the collection axis, the resulting state can be analytically computed: Applying a rotation $R_x(\beta)$ to the atomic dipoles gives

$$R_x(\beta)\,\tilde{\boldsymbol{d}}_{\sigma^+} = \frac{1}{\sqrt{3}} \begin{pmatrix} 1 \\ \mathrm{i}\sin\beta \\ -\mathrm{i}\cos\beta \end{pmatrix}, \qquad R_x(\beta)\,\tilde{\boldsymbol{d}}_{\pi} = \frac{1}{\sqrt{3}} \begin{pmatrix} 0 \\ \cos\beta \\ \sin\beta \end{pmatrix}, \qquad (4.39)$$

and evaluating the corresponding fibre-coupled amplitudes from eq. (4.31) yields the ion–photon state as

$$\left|\psi(\beta)\right\rangle_{IP} = \tilde{I}_0 \left(|0_I H_P\rangle + \mathrm{i}\sin\beta\, |0_I V_P\rangle + \cos\beta\, |1_I V_P\rangle \right). \qquad (4.40)$$

As the magnetic field direction is rotated into the collection axis, the amount of polarisation mixing increases, and the fraction of π emission coupled diminishes. A short calculation [HH99] shows the fully entangled fraction of $\left|\psi(\beta)\right\rangle_{IP}$ to be

$$\mathcal{F}_{\boldsymbol{B}\text{rot}}(\beta) = \frac{1}{2}\left(1 + |\cos\beta|\right), \qquad (4.41)$$

as shown in fig. 4.16, while the fraction of collected photons is constant. $F_{\boldsymbol{B}\text{rot}}(\beta) > 0.995$ for $|\beta| < 8.11°$.

4.3.4   Ion-induced imperfections

Up to now, we have regarded the ion to be a point-like source located at a perfectly known point in space, and the above considerations would apply equally to any other quantum dipole emitter. This is, of course, a rather crude approximation. Among other things, the motional wave function of the ion confined in the three-dimensional harmonic potential would have a non-zero extent even at zero temperature, there is additional motion due to the radio-frequency trapping potential, and the internal electronic state of the ion is also subject to decoherence. We will briefly discuss the effect of these non-idealities in the following.

Secular motion. During the earlier description of the coupling between ion and electromagnetic radiation mode, we have completely neglected atomic motion. This is a good approximation for describing the quantum dynamics of spontaneous emission into vacuum, at least as long as the mode density near the ion position is not made appreciably anisotropic (such as e.g. using an optical cavity). Even if we regard internal state and (classical) motion as completely separate (and still ignore the effect of the photon's momentum, etc.), transforming the resulting field from the atomic frame to the laboratory rest frame
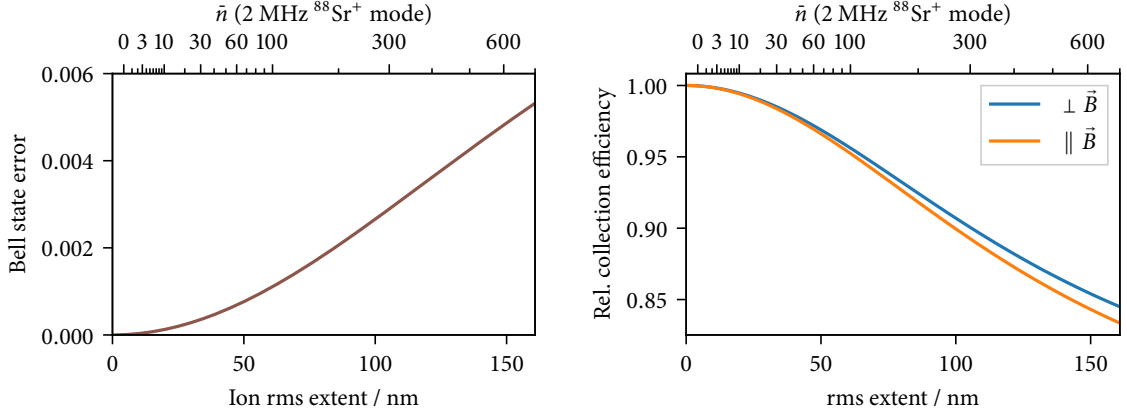
Figure 4.17: Effect of secular ion motion on the fibre-coupled performance, assuming a $2\pi \times 2\,\text{MHz}$ harmonic mode of motion aligned with the given direction perpendicular to the collection axis. Both directions of motion yield the same amount of loss in fidelity; the collection efficiency remains slightly higher for motion perpendicular to the magnetic field as the peak shifted by the $\hat{z}$ dipole component is sampled.

will, for translations along the optical axis, still modulate the phase envelope of the emitted photons, and in either case smear out the apparent source position.

The first effect, phase modulation of the emitted photons, can be highly relevant to two-photon interference under common experimental conditions and is discussed in more detail in §4.5.1. Since the differential effect on the two photon polarisation modes is entirely negligible to the similar wavelength, though, we do not consider it in detail here.

To estimate the strength of effects of the second kind, that is, due to modulation of the ion position, we approximate the positional probability density by a Gaussian, and then consider the ensemble average across the ion–photon states corresponding to the respective emitter positions (the effects of which were discussed earlier). The variance in position of a thermal state (see eq. (2.30)) is

$$(\sigma_x)^2 := \text{tr}(x^2\rho) - \underbrace{(\text{tr}(x\,\rho))^2}_{=0} = x_0{}^2 \frac{1}{1+\bar{n}} \sum_{n=0}^{\infty} \left(\frac{1}{1+\frac{1}{\bar{n}}}\right)^n \underbrace{\langle n|(a+a^\dagger)^2|n\rangle}_{=2n+1} = x_0{}^2(2\bar{n}+1),$$

$$(4.42)$$

giving

$$\sigma_x = \sqrt{x_0\,(2\bar{n}+1)}. \tag{4.43}$$

We have neglected any intrinsic micromotion here, as the secular trap frequencies are typically chosen to be much lower than the Paul trap radio-frequency drive (see next section). A typical trap frequency in ion-trap experiments such as ours might be $\omega_m = 2\pi \times 2\,\text{MHz}$, and the mass of an $^{88}\text{Sr}^+$ ion is $m = 1.46{\times}10^{-25}\,\text{kg}$, yielding the ground-state wave function spread as $x_0 = 5.4\,\text{nm}$ (or 36 nm in the magnified image).

Motion in direction of the axis will only slightly reduce the collection efficiency, as discussed in §4.3.1. The expected performance figures for a normal distribution in ion position across one coordinate axis with a given standard deviation is shown in fig. 4.17.
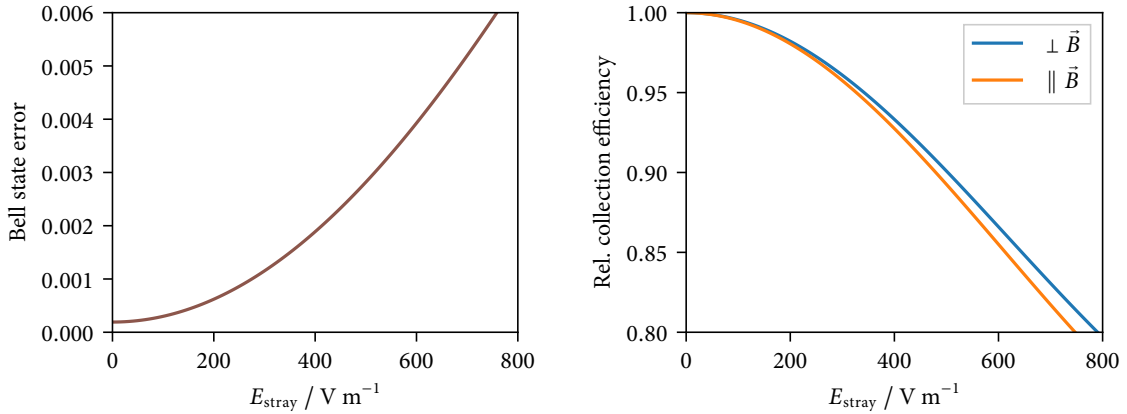
Figure 4.18: Fibre-coupled performance under in-plane micromotion, assuming a 3 MHz radial mode of motion aligned with the given direction perpendicular to the collection axis and 50 MHz trap RF, and $\overline{n} = 10$ for the secular motion.

The reduction in state fidelity is very similar for both transversal directions, as discussed in fig. 4.13. The fidelity remains above 99.5 % for root-mean-square extents below 154 nm, corresponding to $\overline{n} \approx 410$ for the above motional mode parameters. The 95 % collection efficiency threshold is reached at 64 nm, corresponding to $\overline{n} \approx 71$.

Average state occupation numbers of $\overline{n} \approx 10$ are easily reached in typical experiments by Doppler cooling alone, where the expected error is on the 10^{-4}-level. Thus, thermal motion does not appear a significant limitation for ion–photon experiments at present.

Excess micromotion. Radio-frequency Paul traps provide only an approximation to a three-dimensional harmonic potential for charged particles; the effect of the oscillating quadrupole trap fields at frequency ω_{RF} can be reduced to a two-dimensionally confining pseudopotential only in approximation. As discussed in §2.2.1, stray electric fields that shift the ion's equilibrium position off the null of the RF quadrupole potential, lead to modulation of the ion trajectories at the trap drive frequency Ω_{rf}.

The effect of this additional modulation in position on the average fibre-coupling efficiency is shown in fig. 4.18 for micromotion perpendicular to the optical axis, and in fig. 4.19 for micromotion along the optical axis. For in-plane micromotion, the 99.5 % fidelity threshold is only reached at a stray field of 687 V m^{-1}, where the modulation amplitude is already 180 nm. Similarly, the coupling only degrades to 95 % at 331 V m^{-1} for in-plane and 100 V m^{-1} for out-of-plane micromotion. In a typical ion-trap experiment, the stray fields will be compensated to a much better degree than this to avoid unwanted modulation of the laser beams used to manipulate the ion state (that is, keep the micromotion amplitude negligible compared to the ≈ 400 nm wavelength of the laser beams). As such, this positional contribution to ion–photon entanglement errors will usually be negligible.
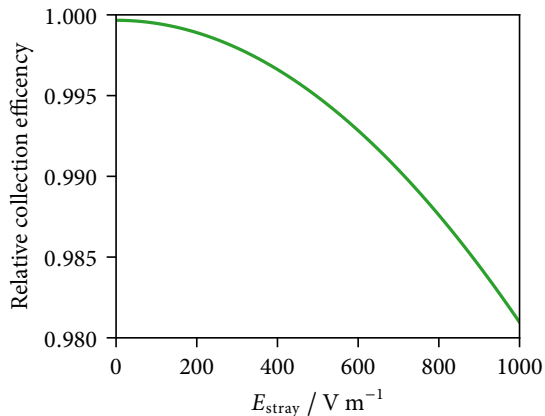
Figure 4.19: Fibre-coupled performance under out-of-plane micromotion, assuming a 3 MHz radial mode of motion parallel to the collection axis with $\bar{n} = 10$, and 50 MHz trap RF.
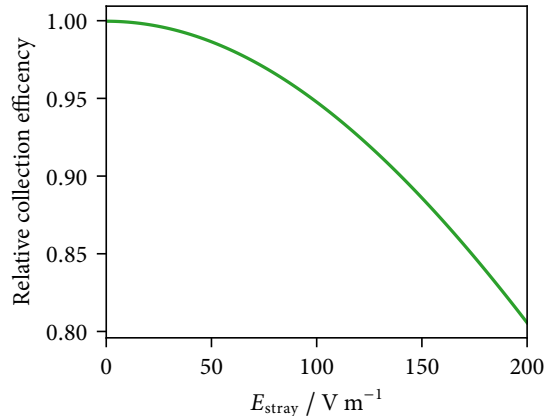
Figure 4.20: Fibre-coupled performance under both defocus and out-of-plane micromotion introduced by a stray electric field of the given strength, assuming a 3 MHz radial mode of motion parallel to the collection axis with $\bar{n} = 10$ for the secular motion, and 50 MHz trap RF.

Note that we assumed that the imaging system was adjusted to compensate for the static shift r_{stray} in the ion position so that just the modulation remains. If this is not the case, e.g. if the stray field varies appreciably over the duration of a given experiment, the positional misalignment introduced by the stray field will affect fidelity and rate as discussed in §4.3.1; shown for the $\hat{z}$ direction in fig. 4.20.

Multiple excitation.   When deriving the ion–photon state following spontaneous decay in section 2.5, we assumed that the ion starts in the excited state with unit probability. For a state with finite lifetime (and a short lifetime is of course necessary to efficiently generate photons through spontaneous emission!) this is only possible to some approximation; the driving field to excite it to that state will need to be applied for some finite duration.

If, during that excitation step, the ion already decays from the excited state, and is excited again, two photons are emitted. If the first photon is detected (which will be most likely), no coherent relationship to the ion state will be observed. This can be modelled exactly, but we can obtain a conservative estimate by considering the admixture of an incoherent mixture of final ion states (according to branching probabilities), with the total weight given by the double-excitation probability.

This is a significant consideration for some experiments where a weak drive is employed [Bli+04], and also for some cavity-assisted schemes [Mer+20; Wal+20]. If a single pulse from a mode-locked laser is used, however, its duration can still be significantly shorter than the excited state lifetime ($\approx 5\,\text{ps}$ in our experiment, compared to the 7.39 ns $P_{1/2}$ state lifetime in $^{88}\text{Sr}^+$). As a result, the intrinsic double excitation probability is negligible, as long as it is not exacerbated due to unintended leakage of laser radiation resonant with the relevant ion transitions after the excitation step (e.g. 422 nm, 1092 nm for $^{88}\text{Sr}^+$).

4.3.5   Other error sources

We briefly mention some other imperfections that will likely be relevant to experimental realisations, but were not explicitly discussed in the above.

Thus far, we have assumed the imaging system to be free of optical aberrations. This is not necessarily a good assumption in practice, as the technical tolerances required to obtain diffraction-limited performance become increasingly challenging to meet for larger numerical apertures. If a parameterisation of the aberrations is known for a particular lens design, it could easily be integrated into the above calculations by multiplying the field in eq. (4.19) by the corresponding pupil function before evaluating the diffraction integral.[10]

There will also be some loss in observed fidelity associated with how the entangled state is used. The ion qubit state will typically be stored for some duration after the emitted photon is observed, leading to decoherence. Limited timing resolution or jitter in the photon detection chain will similarly cause errors in the phase reference for the ion, leading to decoherence as well. The motional state of the ion will typically also induce errors in the coherent operations used to make use of the created state (e.g. tomography for analysis), and qubit state readout will not be perfect.

On the photonic side, we have also neglected any depolarising effects in the collection optics. Given that manipulation of the qubit state would typically happen in fibres or well-collimated free-space beams (where birefringent elements and polarising beamsplitters are available in high quality), the main concern here is potentially non-uniform birefringence across the high-NA collection optics. Slow birefringence drifts in the collection fibres due to changes in mechanical stress or temperature will also limit the ultimate performance of the system, but can be actively corrected if necessary.

4.4   Heralded ion–ion entanglement

Thus far, we have seen that the internal state of the ion after excitation and spontaneous decay is entangled with the polarisation of photons filtered through a single-mode fibre in a way that is maximally entangled conditional on the presence of a photon. We now consider two such nodes, excited at the same time[11], the photons from which are brought together at a central heralding station with the aim of creating a maximally entangled ion–ion state.

Any heralded entanglement generation procedure of this kind is necessarily probabilistic, at least for low photon collection and transmission efficiencies: without knowledge of

---

[10]The performance of high-NA objectives being heavily dependent on their precise alignment, meaningful results are somewhat hard to obtain in practice, however. While it is possible to experimentally fit lower-order models to camera images (see e.g. [Won+16]), a relay imaging system with extra magnification will be necessary for useful results due to the small size of the target fibre mode, which will have to be carefully designed to minimise additional aberrations.

[11]As we shall see, what is actually relevant is that the photon wavepackets arrive at the central station at the same time.

the state of the spontaneous emission photons, the ion states are, with overwhelming probability, just an incoherent mixture of the qubit states, and entanglement is monotonic under local operations and classical control[12].

It is well-known (§2.5) that such a probabilistic entanglement swapping operation, which in theory yields perfect Bell states when a herald signal is asserted, can be implemented in linear optics. This is achieved with a partial Bell basis measurement apparatus based on beamsplitters, where the detection of photons in certain combinations of two output ports heralds the successful (but destructive) projection of the photons onto a Bell state. If the ion–photon input states were maximally entangled, the resulting ion–ion state will thus be maximally entangled too. Of course, this is only exactly true assuming ideal optical components and detectors, as well as photon wavepackets that are perfectly indistinguishable. Aiming to implement such a procedure in practice, especially if the ion–photon input states are close to ideal (as in this work), it is important to understand how sensitive the results are to different practical imperfections.

This is precisely the goal for this section: to derive the ion–ion states corresponding to post-selection on a particular combination of clicks in the single-photon detectors placed at the outputs of a linear-optics Bell basis measurement apparatus, and to do so in a fashion where both non-ideal input states and non-ideal optical devices and components are easily incorporated.

4.4.1   Problem statement

A reduced, schematic representation of the Bell-basis measurement apparatus from fig. 3.7 is shown in fig. 4.21. The inputs are the two photon wavepackets guided by single-mode optical fibres (and the vacuum on the two unused ports of the polarising beamsplitters), the outputs are connected to single-photon detectors with some quantum efficiency $\eta_{\text{det}} \in [0, 1]$, and the components of the apparatus transform the state by interacting with modes in a pairwise fashion. We describe the photonic states using temporal wave-packet modes (see §2.3.1), and do not explicitly consider the geometry or propagation through the apparatus, but just label the modes by a spatial mode ("beam") index $s \in \{A, B, C, D\}$ and their polarisation Π (for which we will here explicitly use the horizontal and vertical directions relative to the optical breadboard to which all the components are mounted, $\Pi \in \{H, V\}$). Let $\rho^s_{I^s P}$ be the ideal ion–photon state[13] at each input of the heralding station,

$$\rho^s_{I^s P} = \eta_{422} \eta^s_{\text{fibre}} \, |\psi^s\rangle\langle\psi^s| + \left(1 - \eta_{422} \eta^s_{\text{fibre}}\right) \rho^s_{\text{dark},I^s} \otimes |0\rangle\langle 0| \,, \qquad (4.44)$$

---

[12]A note in passing: The Bell-state fidelity or fully entangled fraction $\mathcal{F}$ is not itself a LOCC monotone; cf. [Bad+00]. For the present discussion, where the aim is not to e.g. demonstrate the abstract (in)feasibility of a certain scheme, but to analyse the performance of a particular implementation, this is not of concern; readily computed LOCC monotones such as the logarithmic negativity [Ple05] exist, however.

[13]In $\rho^s_{I^s P}$ and throughout, the subscript s is not an exponent but labels the system that the state or parameter describes. I^A denotes the ion in node Alice, I^B that in Bob, and P the general multi-mode photonic system.
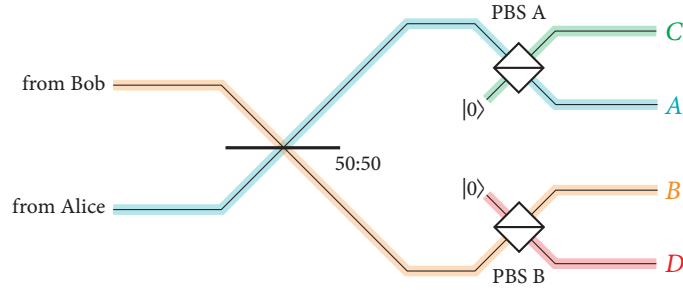
Figure 4.21: Schematic representation of the partial Bell-basis measurement apparatus, with mode labels as used throughout this section. The inputs are single-mode fibres from the ion trap nodes Alice and Bob, the output modes A to D are again coupled into single-mode fibres to single-photon detectors (see fig. 3.7 for a sketch of the actual experimental implementation). The colouring suggests the spatial mode naming convention used throughout this section: rather than labelling e.g. the input and output modes of a beam-splitter with different letters, we consider each component to cause pairs of modes to interact (as expressed by a unitary transformation on the associated creation operators). The modes retain the names they would have had if the component were not present. Modes A and B perfectly overlap on the $50:50$ beamsplitter, that is, the spatial modes for photons from Alice and from Bob are indistinguishable.

where η_{422} is the fraction of $P_{1/2}$ decays on the desired 422 nm transition (see §4.1) and η_{fibre} is the efficiency of collection into and transmission through the fibre. ρ^s_{dark,I^s} is the ion state when the photon has not been transmitted by the fibre and has been traced out; for efficient excitation[14] it is

$$\rho^s_{\text{dark},I^s} = \frac{1}{1 - \eta_{422}\eta^s_{\text{fibre}}} \Bigg((1 - \eta_{422}) \rho_{1092,I^s} + \eta_{422}\left(\left(\frac{2}{3} - \frac{\eta^s_{\text{fibre}}}{2}\right)|0\rangle\langle 0| + \left(\frac{1}{3} - \frac{\eta^s_{\text{fibre}}}{2}\right)|1\rangle\langle 1|\right)\Bigg), \quad (4.45)$$

but as the protocol considered here is based on the detection of one photon from each ion, its form is not relevant for the present discussion. $|\psi^s\rangle$ is the maximally entangled state between ion state and photon wavepackets,

$$|\psi^s\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle \otimes \int_{-\infty}^{\infty} \phi_0(t - \Delta t^s)\, b^{\dagger}_{sH}(t)\, \mathrm{d}t + |1\rangle \otimes \int_{-\infty}^{\infty} \phi_1(t - \Delta t^s)\, b^{\dagger}_{sV}(t)\, \mathrm{d}t\right)|0\rangle, \quad (4.46)$$

where we have taken Δt^s to be the propagation delay from ion to the entangler input ($\Delta t := L/c$ for a free-space distance L, appropriately weighted by the refractive index in fibres). We have assumed that the polarisations corresponding to σ^+ and π decays are exactly aligned with the H and V axes for clarity of notation; we will consider the effect of residual fibre birefringence in §4.5.2.

For this analysis, however, we do not wish to assume the ion–photon state to be ideal, but want to be able to propagate any number of imperfections, such as those discussed

---

[14]If the picosecond laser excitation from $|0\rangle = \left|5s\,^2S_{1/2}, m_J = -\frac{1}{2}\right\rangle$ to $|e\rangle = \left|5p\,^2P_{1/2}, m_J = \frac{1}{2}\right\rangle$ is not perfectly efficient, the ratio of $2:1$ in the lower state populations would be biased further towards $|0\rangle$.

in §4.3, though the Bell-state measurement apparatus to model their effect on the final states. We thus consider the input from each system, Alice ($s = A$) or Bob ($s = B$), to be given by a density operator. To keep the notation concise in the rest of the chapter, we anticipate the fact that we will only consider events heralded by the presence of two photons and disregard the $\rho^s_{\mathrm{dark},I^s}$ parts, instead working with sub-normalised density operators $\rho^s_{I^sP}$ which only have support on the one-photon subspace, their trace given by the total system collection efficiency η^s. η^s is a product of η_{422}, η^s_{fibre}, and the efficiencies of any other imperfect state preparation, excitation or transmission processes; that is, the probability to actually find a photon at the output of the fibre from s following the initial excitation stage. We can express such a state generically "in components" as

$$\rho^s_{I^sP} = \eta^s \sum_{i,j\in\{0,1\}} |i\rangle\langle j| \otimes \tau^s_{ij}, \quad \tau^s_{ij} \coloneqq \sum_{o,p\in\{H,V\}} \int\limits_{-\infty}^{\infty} \int\limits_{-\infty}^{\infty} \rho^s_{ijop}(l,r)\, b^\dagger_{no}(l)\, |0\rangle\langle 0|\, b_{np}(r)\, \mathrm{d}r\, \mathrm{d}l. \quad (4.47)$$

where the $\rho^s_{ijop}(l,r) \in \mathbb{C}$ give the coefficients in this basis, with discrete ion state indices i,j, polarisation indices o,p, and the continuous time parameters l,r. The "photonic part" τ^s_{ij} of the tensor product depends on the ion state indices i,j unless ion and photon are in a product state. For the ideal input state from eq. (4.46), $\rho^s_{ijop}(l,r) = \frac{1}{2}\delta_{\pi(i)o}\delta_{\pi(j)p}\phi_i(l - \Delta t^s)\overline{\phi_j(r - \Delta t^s)}$, where π specifies the ion state/polarisation mapping $\pi(0) = H$, $\pi(1) = V$.

We are interested in the density operator describing the joint state of the two ions after a herald signal given by, in this case, two clicks in the single-photon detectors on the output of the Bell-state measurement apparatus. The beamsplitter network is usefully[15] described by a unitary operator U_P acting on the photonic part of the system, such that the state at the output of the network is given by $\rho^{\mathrm{out}} \coloneqq (\mathbb{1}_{I^AI^B} \otimes U_P)(\rho^A_{I^AP} \otimes \rho^B_{I^BP})(\mathbb{1}_{I^AI^B} \otimes U^\dagger_P)$.

As discussed in §2.3.1, the detection of two photons in modes $s\Pi$ and $s'\Pi'$ at times t and t' is a projective measurement given by a measurement operator $M \coloneqq \langle 0| b_{s\Pi}(t)\, b_{s'\Pi'}(t')$ and projector $M^\dagger M$, or for detectors with a finite efficiency[16] η_{det}, the POVM element $(\eta_{\mathrm{det}})^2 M^\dagger M$. Detection of photons in a given time window is expressed by an appropriate integration over the time variables. Given a POVM element F, the sub-normalised ion–ion state following this measurement is

$$\begin{aligned}
(\rho_{\wedge F})_{I^AI^B} &= \mathrm{tr}_P\left((\mathbb{1}_{I^AI^B} \otimes F)\, \rho^{\mathrm{out}}\right) \\
&= \mathrm{tr}_P\left((\mathbb{1}_{I^AI^B} \otimes F)(\mathbb{1}_{I^AI^B} \otimes U_P)(\rho^A_{I^AP} \otimes \rho^B_{I^BP})(\mathbb{1}_{I^AI^B} \otimes U^\dagger_P)\right) \\
&= \eta^A\eta^B \sum_{i^Aj^Ai^Bj^B} |i^A\rangle\langle j^A| \otimes |i^B\rangle\langle j^B| \otimes \mathrm{tr}\left(F U_P\left(\tau^A_{i^Aj^A} \otimes \tau^B_{i^Bj^B}\right)U^\dagger_P\right).
\end{aligned} \quad (4.48)$$

---

[15]The unitarity of the beamsplitter network is technically only an approximation; it consists only of linear optics, but is not photon-number preserving due to small amounts of loss. However, as this loss is approximately uniform and we are only interested in the case post-selected on the presence of a pair of photons here, it can be easily accounted for in η^s or η_{det}. If necessary, the conjugation by U_P could easily be replaced by the application of an arbitrary map either way.

[16]For conciseness, we assume the detection efficiency to be the same across all detectors here, as it merely results in a prefactor to the overall probability-weighted density matrix. Individual detector efficiencies can easily be reinstated by replacing the prefactor by an appropriate product of efficiencies.

The trace of $(\rho_{\wedge F})_{I^A I^B}$ gives the probability of observing two photons as specified by F; if F consists directly of temporal mode operators $b_{s\Pi}(t)$ rather than their integral, it has the meaning of a probability density for the detection of a photon at t.

4.4.2   Evolution and post-selected states in the ideal case

To describe the unitary evolution, we will consider the apparatus in a "lumped-element" model, concatenating the effect of the various optical elements while neglecting the free-space paths between them (which are inconsequential here[17]). Each element couples excitations in two modes in a way that can be described by transforming their creation operators $b^\dagger_{s\Pi}$ and $b^\dagger_{s'\Pi'}$ into a new pair of creation operators $\tilde{b}^\dagger_{s\Pi}$ and $\tilde{b}^\dagger_{s'\Pi'}$ describing excitations in the output modes according to the corresponding two-level unitary matrix $U \in \mathrm{U}(2)$. We will write $\mathcal{V}_{x\leftrightarrow y}(U)$ for this action of U on the modes x and y. This could be evaluated by symbolically applying the respective transformations to the input state in eq. (4.48), though this is cumbersome in the general case; appendix A describes a way to, taking advantage of the fact that the system never contains more than a few (here, two) excitations at the same time, evaluate these expressions through matrix multiplication in a truncated multi-mode Fock space.

In the absence of any birefringence or other imperfections, the total evolution operator is simply given by a series of beamsplitters:

$$\begin{aligned}
U_P = {} & \mathcal{V}_{BV\leftrightarrow DV}\left(U_{\mathrm{BS}}(\theta_V^{\mathrm{PBS\,B}})\right) \cdot \mathcal{V}_{BH\leftrightarrow DH}\left(U_{\mathrm{BS}}(\theta_H^{\mathrm{PBS\,B}})\right) \cdot \\
& \mathcal{V}_{AV\leftrightarrow CV}\left(U_{\mathrm{BS}}(\theta_V^{\mathrm{PBS\,A}})\right) \cdot \mathcal{V}_{AH\leftrightarrow CH}\left(U_{\mathrm{BS}}(\theta_H^{\mathrm{PBS\,A}})\right) \cdot \\
& \mathcal{V}_{AV\leftrightarrow BV}\left(U_{\mathrm{BS}}(\theta_V^{50:50})\right) \cdot \mathcal{V}_{AH\leftrightarrow BH}\left(U_{\mathrm{BS}}(\theta_H^{50:50})\right).
\end{aligned}$$

(4.49)

$U_{\mathrm{BS}}(\theta)$ is the unitary matrix for a beamsplitter with reflectance $(\sin\theta)^2$, as discussed in §2.3.2. In the ideal case, the non-polarising beamsplitter has a $50:50$ splitting ratio independent of polarisation, $\theta_H^{50:50} = \theta_V^{50:50} = \frac{\pi}{4}$, and the polarising beamsplitters are perfectly transmissive for H polarisation, perfectly reflective for V polarisation: $\theta_H^{\mathrm{PBS\,A}} = \theta_H^{\mathrm{PBS\,B}} = 0$, $\theta_V^{\mathrm{PBS\,A}} = \theta_V^{\mathrm{PBS\,B}} = \frac{\pi}{2}$. We will write U_P^{ideal} for eq. (4.49) with these parameters.

At the output, each mode is coupled into a polarisation-insensitive detector[18]. We hence label the detection events using the spatial mode names (see table 4.1 for the correspondence to the detector indices used in the experimental chapters), for which the corresponding POVMs are sums over the individual polarisations. With perfect PBSes, the

---

[17]As long as the photon wavepackets are timing-aligned on the non-polarising $50:50$ beamsplitter (fig. 4.21), the exact amount of time delay in the various paths before or after the beamsplitter only contributes an uninteresting phase shift – we do not use a scheme such as homodyne detection that would be sensitive to the absolute optical phase. If the apparatus contained paths that split and recombine, as e.g. in a Mach–Zehnder interferometer, we would need to track the temporal evolution even without time/phase resolution on the detectors to capture the differential phase shift resulting from path length differences.

[18]Just as noted regarding any differences in efficiency between the detectors earlier, a dependence on polarisation would be easy to insert – as we only consider the state post-selected on the presence of the photons in the experiment, this would only modify the probability of observing the various heralds accordingly.

Spatial mode notation	A	B	C	D
Experimental apparatus	APD3	APD0	APD2	APD1

Table 4.1: Mapping between spatial mode names used here and photodetector names used in the experimental chapters.

output modes AV, BV, CH, DH will of course not be populated; in fact, we could have directly assigned the modes AH, AV, BH, BV after the $50 : 50$ beamsplitter to the four detectors. Introducing the extra modes will allow us to easily introduce imperfections in the next section, though.

Now, we can finally compute the ion–ion density matrices $\rho_{\wedge x \wedge y}(t, t')$ corresponding to the detection of photons in modes x and y at times t and t', respectively (we will drop the ion system subscript $I^A I^B$ from here on out). By convention, we will arrange the modes x and y in lexicographic order and match them up with the time variables accordingly. The matrices are sub-normalised; their trace gives the joint probability density for detecting two photons at the given times. For the four combinations of detectors of opposite polarisations, we obtain (for $U_P = U_P^{\text{ideal}}$)

$$
\rho_{\wedge A \wedge C}(t, t') = \rho_{\wedge B \wedge D}(t, t') = \frac{\eta^A \eta^B (\eta_{\text{det}})^2}{8} \Gamma^2 e^{-\Gamma t} e^{-\Gamma t'} \, |\Psi^+\rangle\langle\Psi^+| ,
$$
$$
\rho_{\wedge A \wedge D}(t, t') = \rho_{\wedge B \wedge C}(t, t') = \frac{\eta^A \eta^B (\eta_{\text{det}})^2}{8} \Gamma^2 e^{-\Gamma t} e^{-\Gamma t'} \, |\Psi^-\rangle\langle\Psi^-| .
$$

$$(4.50)$$

In the ideal case, we thus obtain a maximally entangled state on detection of a two-photon coincidence, independent of the photon detection time. Integrating over the entire exponential tail, each click pattern occurs with probability $\frac{1}{8}\eta^A \eta^B (\eta_{\text{det}})^2$, for an overall probability of $\frac{1}{2}\eta^A \eta^B (\eta_{\text{det}})^2$ of a successful herald.

The same-polarisation coincidences $A \wedge B$ and $C \wedge D$ do not occur at all in the perfect case; this is the Hong–Ou–Mandel effect [HOM87]. Both these and the cases of two photons in the same mode, which we cannot detect in our experiment, would just lead to non-entangled states $|00\rangle\langle00|$ or $|11\rangle\langle11|$ – if two photons of the same polarisation are detected, there can be no "which-path" ambiguity about which ion is in what state.

4.5   Ion–ion entanglement imperfections

One obvious source of deviations of the post-herald ion–ion state from the ideal maximally-entangled state are those issues also affecting the ion–photon-entangled states it resulted from; we have already discussed those in §4.3. Here, we will catalogue a number of additional effects affecting the ion–ion entanglement beyond those, discussing their effect on the remote entanglement quality in isolation to provide guidelines for the selection of optical components, alignment tolerances, etc. (though the model could of

course equally be used to predict the outcome from any number of these issues present at the same time).

4.5.1   Ideal entangler, non-identical photon wavepackets

For the first set of imperfections, we will consider those where the entanglement swapping apparatus is perfect, but the temporal input modes are not perfectly identical. Now writing ϕ_0^A, ϕ_1^A for the amplitudes from Alice and ϕ_0^B, ϕ_1^B for those from Bob, we can calculate the density matrices for the four heralding patterns in the same way as for eq. (4.50), and obtain

$$
\rho_{\wedge A \wedge C}(t,t') = \rho_{\wedge B \wedge D}(t,t') =
$$

$$
\frac{\eta^A \eta^B (\eta_{\text{det}})^2}{16}
\begin{pmatrix}
0 & 0 & 0 & 0 \\
0 & |\phi_0^A(t)|^2 |\phi_1^B(t')|^2 & \phi_0^A(t)\phi_1^B(t')\overline{\phi_1^A(t')\phi_0^B(t)} & 0 \\
0 & \phi_1^A(t')\phi_0^B(t)\overline{\phi_0^A(t)\phi_1^B(t')} & |\phi_1^A(t')|^2 |\phi_0^B(t)|^2 & 0 \\
0 & 0 & 0 & 0
\end{pmatrix},
$$

$$
\rho_{\wedge A \wedge D}(t,t') = \rho_{\wedge B \wedge C}(t,t') =
$$

$$
\frac{\eta^A \eta^B (\eta_{\text{det}})^2}{16}
\begin{pmatrix}
0 & 0 & 0 & 0 \\
0 & |\phi_0^A(t)|^2 |\phi_1^B(t')|^2 & -\phi_0^A(t)\phi_1^B(t')\overline{\phi_1^A(t')\phi_0^B(t)} & 0 \\
0 & -\phi_1^A(t')\phi_0^B(t)\overline{\phi_0^A(t)\phi_1^B(t')} & |\phi_1^A(t')|^2 |\phi_0^B(t)|^2 & 0 \\
0 & 0 & 0 & 0
\end{pmatrix}.
$$

$$(4.51)$$

For non-identical envelopes, the Hong–Ou–Mandel-suppressed same-polarisation patterns are no longer entirely forbidden. The density matrices are

$$
\rho_{\wedge A \wedge B}(t,t') = \frac{\eta^A \eta^B (\eta_{\text{det}})^2}{16} |\phi_0^A(t)\phi_0^B(t') - \phi_0^A(t')\phi_0^B(t)|^2 \, |0\rangle\langle 0| ,
$$

$$
\rho_{\wedge C \wedge D}(t,t') = \frac{\eta^A \eta^B (\eta_{\text{det}})^2}{16} |\phi_1^A(t)\phi_1^B(t') - \phi_1^A(t')\phi_1^B(t)|^2 \, |1\rangle\langle 1| .
$$

$$(4.52)$$

This shows that equal-time coincidences remain entirely forbidden, $\operatorname{tr}\rho_{\wedge A \wedge B}(t,t) = \operatorname{tr}\rho_{\wedge C \wedge D}(t,t) = 0$ for all t, but coincidences at different times $t \neq t'$ can be allowed if the wavepackets for the respective polarisation are not identical. This matches earlier results on time-resolved Hong–Ou–Mandel interference in a purely photonic context [Leg+03; Leg+06]. We will see a concrete example for the case of a difference in magnetic fields later.

From eq. (4.50), we saw that in the ideal case, the ion–ion state does not depend on the detection times. Furthermore, to optimise the heralding probability, clicks at all detection times should be allowed. This is indeed the mode we typically operate the experiment in, where we only use the photon detector timestamps to define a broad coincidence window to exclude counts before the excitation pulse, and after ~ 4 decay time constants have

elapsed as a fairly arbitrary cut-off to reduce the influence of dark counts. In this case, the states are (in the limit where the exponential tail is not cut off at all) given by

$$\rho_{\wedge A \wedge C} = \rho_{\wedge B \wedge D} := \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \rho_{\wedge A \wedge C}(t, t') \, dt' dt =$$

$$\frac{\eta^A \eta^B (\eta_{\text{det}})^2}{16} \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & \|\phi_0^A\|^2 \|\phi_1^B\|^2 & \langle \phi_0^B, \phi_0^A \rangle \langle \phi_1^A, \phi_1^B \rangle & 0 \\ 0 & \langle \phi_0^A, \phi_0^B \rangle \langle \phi_1^B, \phi_1^A \rangle & \|\phi_1^A\|^2 \|\phi_0^B\|^2 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}$$

$$\rho_{\wedge A \wedge D} = \rho_{\wedge B \wedge C} := \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \rho_{\wedge A \wedge D}(t, t') \, dt' dt =$$

$$\frac{\eta^A \eta^B (\eta_{\text{det}})^2}{16} \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & \|\phi_0^A\|^2 \|\phi_1^B\|^2 & -\langle \phi_0^B, \phi_0^A \rangle \langle \phi_1^A, \phi_1^B \rangle & 0 \\ 0 & -\langle \phi_0^A, \phi_0^B \rangle \langle \phi_1^B, \phi_1^A \rangle & \|\phi_1^A\|^2 \|\phi_0^B\|^2 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad (4.53)$$

where the diagonal elements are 1 as the envelope functions are normalised as $\|\phi_0^A\| = \|\phi_0^B\| = \|\phi_1^A\| = \|\phi_1^B\| = 1$, and for the Hong–Ou–Mandel-forbidden patterns

$$\rho_{\wedge A \wedge B} := \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \rho_{\wedge A \wedge B}(t, t') \, dt' dt = \frac{\eta^A \eta^B (\eta_{\text{det}})^2}{8} \left(1 - |\langle \phi_0^B, \phi_0^A \rangle|^2\right) |0\rangle\langle 0|,$$

$$\rho_{\wedge C \wedge D} := \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \rho_{\wedge C \wedge D}(t, t') \, dt' dt = \frac{\eta^A \eta^B (\eta_{\text{det}})^2}{8} \left(1 - |\langle \phi_1^B, \phi_1^A \rangle|^2\right) |1\rangle\langle 1|. \quad (4.54)$$

For the valid patterns, $\rho \in \{\rho_{\wedge A \wedge C}, \rho_{\wedge A \wedge D}, \rho_{\wedge B \wedge C}, \rho_{\wedge B \wedge D}\}$, $\operatorname{tr} \rho = \eta^A \eta^B (\eta_{\text{det}})^2/8$, and, as ρ is of the form (2.14),

$$\mathcal{F}\left(\frac{\rho}{\operatorname{tr} \rho}\right) = \frac{1}{2} \left(1 + |\langle \phi_0^B, \phi_0^A \rangle \langle \phi_1^A, \phi_1^B \rangle|\right). \quad (4.55)$$

If we define the probability of observing the Hong–Ou–Mandel-suppressed patterns relative to the allowed patterns according to eq. (4.54) as $\varepsilon_H := 1 - |\langle \phi_0^B, \phi_0^A \rangle|^2$, $\varepsilon_V := 1 - |\langle \phi_1^B, \phi_1^A \rangle|^2$, then

$$\mathcal{F}\left(\frac{\rho}{\operatorname{tr} \rho}\right) = \frac{1}{2} \left(1 + \sqrt{(1 - \varepsilon_H)(1 - \varepsilon_V)}\right), \quad (4.56)$$

or, in terms of the Bell-state error,

$$\mathcal{E}\left(\frac{\rho}{\operatorname{tr} \rho}\right) = \frac{1}{2} \left(1 - \sqrt{(1 - \varepsilon_H)(1 - \varepsilon_V)}\right) \approx \frac{\varepsilon_H + \varepsilon_V}{4}. \quad (4.57)$$

In other words, if (but only if) the Bell-state measurement apparatus is perfect, then the amount of infidelity resulting from non-identical wavepackets can easily be diagnosed through the probability for Hong–Ou–Mandel-forbidden coincidences to occur.

Temporal misalignment. If the arrival times of the photons from Alice and Bob at the central beamsplitter do not match, the wave-packet overlap decreases accordingly: Without loss of generality, assume that Bob's photons are shifted back by $\Delta\tau \geq 0$, such that $\phi_i^B(t) = \phi_i^A(t - \Delta\tau)$. Then,

$$|\langle\phi_0^B, \phi_0^A\rangle| = |\langle\phi_1^A, \phi_1^B\rangle| = \Gamma \int_{\Delta\tau}^{\infty} e^{-\Gamma\left(t - \frac{\Delta\tau}{2}\right)} dt = e^{-\frac{\Gamma}{2}\Delta\tau}, \tag{4.58}$$

so the Bell-state error from a temporal misalignment is

$$\mathcal{E}_{\Delta\tau} = \frac{1 - e^{-\Gamma|\Delta\tau|}}{2} \approx \frac{\Gamma|\Delta\tau|}{2} + O\left((\Gamma\Delta\tau)^2\right), \tag{4.59}$$

where the absolute value restores validity for general $\Delta\tau$ (as can be seen from the same consideration with the roles of Alice and Bob swapped). For $\mathcal{E}_{\Delta\tau} < 1\,\%$ in $^{88}\mathrm{Sr}^+$, $\Delta\tau <$ 0.15 ns, i.e. up to 4.5 cm in free-space path mismatch can be tolerated.

Note, however, that all the error comes from heralds where one of the photons is detected at $t \in [0, \Delta\tau]$. For all later times, the exponential decay envelope is self-similar and eq. (4.51) yields maximally entangled states. If the temporal alignment in the experiment cannot be established to better than $|\Delta\tau| < \tau$ for a given bound τ, the reduction in fidelity can thus be avoided by only considering detector clicks in $[\tau, \infty)$ relative to the first registered photons to be valid. Thus, temporal misalignment can be regarded as instead having a coincidence rate cost of $e^{-\Gamma\tau}$.

Magnetic field mismatch. Recall from eq. (4.3) that the photon wavepackets are given by $\phi_i(t) := \sqrt{\Gamma}e^{-i\delta_i t} e^{-\frac{\Gamma}{2}t} \Theta(t)$, where $\delta_i = \nu_{ei}B$ is the difference in atomic splitting from the zero-field transition frequency we chose as the reference, as given by the overall Zeeman sensitivity ν_{ei} to the static ambient magnetic field B (in angular units). If the magnetic fields at the two systems differ by $\Delta B := B^A - B^B$, this of course does not individually cause any drop in the ion–photon entanglement fidelity for each system, as long as the qubit frequencies are properly calibrated. In the ion–ion case, however, the wavepacket overlap integrals are

$$\langle\phi_0^B, \phi_0^A\rangle = \Gamma \int_0^{\infty} e^{-\Gamma t} e^{-i\nu_{e0}\Delta B t} dt = \frac{\Gamma}{\Gamma + i\nu_{e0}\Delta B},$$
$$\langle\phi_1^A, \phi_1^B\rangle = \Gamma \int_0^{\infty} e^{-\Gamma t} e^{i\nu_{e1}\Delta B t} dt = \frac{\Gamma}{\Gamma - i\nu_{e1}\Delta B}, \tag{4.60}$$

giving rise to an error

$$\mathcal{E}_{\Delta B} = \frac{1}{2}\left(1 - \frac{1}{\sqrt{\left(1 + \frac{\nu_{e0}^2\Delta B^2}{\Gamma^2}\right)\left(1 + \frac{\nu_{e1}^2\Delta B^2}{\Gamma^2}\right)}}\right) \approx \frac{\nu_{e0}^2 + \nu_{e1}^2}{4\Gamma^2}\Delta B^2 + O(\Delta B^4). \tag{4.61}$$

For the error in ^{88}Sr$^+$ to accumulate to 1 % across the short 7.39 ns photons, a magnetic field mismatch on the 0.2 mT-level would be needed; here, this is almost half the nominal quantisation field.

In other situations where significant qubit frequency mismatch might be more common (e.g. because the frequency cannot be as freely tuned as through the B field here), it can be interesting not to discard the time information. Assume, more generally, that the photonic envelopes from the two ions are related by a pure phase term, $\phi_i(t) \coloneqq \phi_i^A(t) = e^{-i\vartheta_i(t)}\phi_i^B(t)$ for some $\vartheta_i : \mathbb{R} \to \mathbb{R}$. The density matrices from eq. (4.51) are then (without applying any knowledge about the form of $\phi_i(t)$)

$$\rho_{\wedge A \wedge C}(t, t') = \rho_{\wedge B \wedge D}(t, t') =$$

$$\frac{\eta^A \eta^B (\eta_{\mathrm{det}})^2}{16} |\phi_0(t)|^2 |\phi_1(t')|^2 \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & e^{i(\vartheta_1(t')-\vartheta_0(t))} & 0 \\ 0 & e^{-i(\vartheta_1(t')-\vartheta_0(t))} & 1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \qquad (4.62)$$

and with minus signs on the anti-diagonal elements for $\rho_{\wedge A \wedge D}(t, t') = \rho_{\wedge B \wedge C}(t, t')$. For the Hong–Ou–Mandel-suppressed coincidences,

$$\rho_{\wedge A \wedge B}(t, t') = \frac{\eta^A \eta^B (\eta_{\mathrm{det}})^2}{8} |\phi_0(t)|^2 |\phi_0(t')|^2 \left(1 - \cos\left(\vartheta_0(t') - \vartheta_0(t)\right)\right) |0\rangle\langle 0|,$$

$$\rho_{\wedge C \wedge D}(t, t') = \frac{\eta^A \eta^B (\eta_{\mathrm{det}})^2}{8} |\phi_1(t)|^2 |\phi_1(t')|^2 \left(1 - \cos\left(\vartheta_1(t') - \vartheta_1(t)\right)\right) |1\rangle\langle 1|. \qquad (4.63)$$

From eq. (4.62), it is clear that for every particular pair of detection times t and t', the ions are in a pure, maximally entangled state

$$\rho_{|(A \wedge C)}(t, t') = \rho_{|(B \wedge D)}(t, t') = |\psi_{t,t'}\rangle\langle\psi_{t,t'}|, \text{ with } |\psi_{t,t'}\rangle = \frac{1}{\sqrt{2}}\left(|01\rangle + e^{i(\vartheta_0(t)-\vartheta_1(t'))}|10\rangle\right), \qquad (4.64)$$

with a minus sign for the $A \wedge D$ and $B \wedge C$ patterns. If the form of the time dependence is known, an appropriate phase gate (local z-axis rotation) can thus be applied to the ion qubits afterwards to correct for it.

Retuning to the particular case of a magnetic field mismatch, the time dependence is simply $\vartheta_i(t) = \nu_{ei}\Delta B\, t$ giving the phase correction[19] to apply as $e^{i\Delta B(\nu_{e0}t-\nu_{e1}t')}$. This matches, and generalises, the results from an experimental demonstration [Vit+14], where both polarisations were assumed to have the same frequency difference. A phase feed-forward scheme of this kind has also been demonstrated for entanglement swapping in a

---

[19]Recall from §4.1 that we are, for each ion, working in a rotating frame exactly matching the Zeeman qubit splitting, and referenced to the laboratory frame at the time of the excitation pulse. In eliding the temporal evolution through the photon collection optics, optical fibres, and the entangler, we have here chosen $t = t' = 0$ to correspond to the initial edge of the wavepacket, where spontaneous emission photons begin to arrive at the detectors in question.
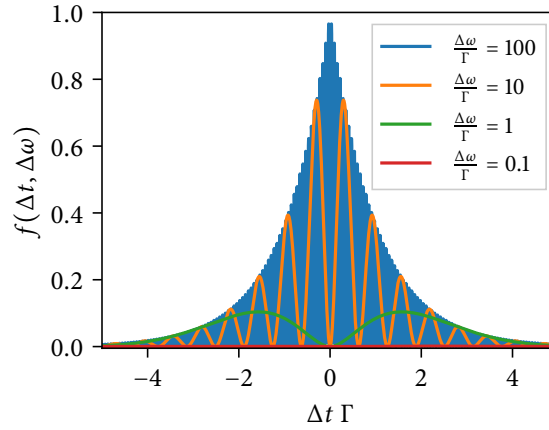
Figure 4.22: Hong–Ou–Mandel suppression of same-polarisation coincidences on opposite sides of the 50:50 beamsplitter if the photon centre frequencies do not match exactly. Shown is the probability density for detecting two photons with a time difference Δt normalised by the linewidth Γ (see eq. (4.65)). As the frequency difference $\Delta\omega$ becomes much larger than the linewidth, the behaviour approaches that of the fully distinguishable case.

purely photonic experiment [Zha+14]. To illustrate the effect on the Hong–Ou–Mandel-suppressed coincidences, let $t' = t + \Delta t$ and average across all detection times t of the other photon:

$$\int_{-\infty}^{\infty} \rho_{\wedge A \wedge B}(t, t + \Delta t)\mathrm{d}t = \frac{\eta^A \eta^B (\eta_{\mathrm{det}})^2}{8}\Gamma f(\Delta t, \nu_{e0}\Delta B)\,|0\rangle\langle 0|\,,$$

$$\text{with } f(\Delta t, \Delta\omega) \coloneqq \frac{1}{2}\mathrm{e}^{-\Gamma|\Delta t|}\left(1 - \cos\left(\Delta\omega\Delta t\right)\right).$$

(4.65)

(same for $C \wedge D$ with $\Delta\omega = \nu_{e1}\Delta B$). The behaviour of the normalised coincidence probability function f is shown in fig. 4.22: While coincidences at $\Delta t = 0$ remain disallowed, the oscillations within the exponential envelope given by the atomic linewidth increase in frequency with $\Delta\omega$, approaching the case of two distinguishable photons for $\Delta\omega \to \infty$.

Periodic phase modulation (e.g. micromotion). A key fact regarding the relative importance of various imperfections is that any change in the optical path lengths, whether through motion of the ion emitters or mechanical changes in the optical system, will be on a similar scale to the optical wavelength (422 nm), which is much smaller than the spatial extent of the photon wavepackets (2 m) or the wavelength corresponding to the Zeeman qubit frequency splitting (21 m). Any such change, whether in the form of slow drifts across many entanglement events, or fast motion across a single photon wavepacket, can thus in excellent approximation be described by the same purely imaginary phase envelope applied to both $\phi_0(t)$ and $\phi_1(t)$.

A well-known feature of two-photon heralding schemes is that they are not sensitive to constant offsets in the absolute optical phase (cf. eq. (4.62) with $\vartheta_0(t) = \vartheta_1(t) = \delta\varphi$). This is a big advantage, as it means that no long-term interferometric stability is required.
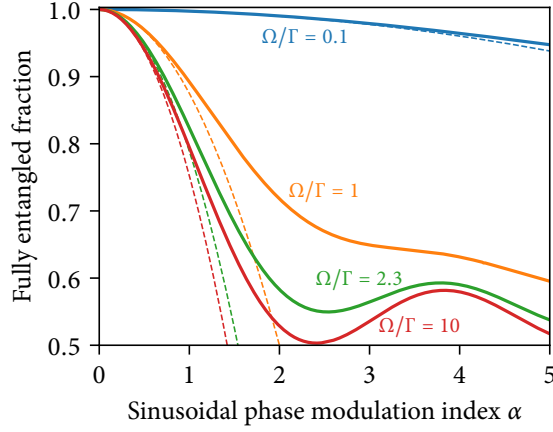
Figure 4.23: Loss in ion–ion Bell state fidelity from a sinusoidal phase modulation $e^{-i\alpha \sin(\Omega t + \varphi)}$ (with uniformly distributed phase φ). The strength of the effect depends on the ratio of modulation frequency to excited state linewidth; an approximately constant phase offset from modulation slow compared to the length of the photons does not have any effect on the ion–ion state. The dashed curves give the approximation $\frac{\Omega^2}{\Gamma^2 + \Omega^2} \left(\frac{\alpha}{2}\right)^2$ (to second order for α small).

However, this does not hold for modulation at time scales comparable to or faster than the linewidth Γ. A prominent source of such modulation in Paul traps is excess micromotion at the trap drive frequency Ω_{rf} (see §2.2.1). We already considered the effects of micromotion in the imaging direction in §4.3.4, and saw that its effect on the coupling efficiency was only mild. Continuing to neglect any higher-order effects on the spontaneous emission process, we consider the effect of a positional modulation $r \sin(\Omega t + \varphi)$ to be a pure phase modulation $e^{-i\alpha \sin(\Omega t + \varphi)}$, where $\alpha := k_n r$ is the modulation index given the angular wavenumber k_n corresponding to the optical transition frequency (which is to good approximation the same for both decay channels). For excess micromotion, $\Omega = \Omega_{\text{rf}}$, and $\varphi = \varphi_{\text{rf}}$ is the phase of the radio-frequency drive with respect to the excitation pulse timing, which varies quickly and, in effect, randomly from entanglement attempt to entanglement attempt. Hence, we consider the ion–photon state of Alice to be in a statistical mixture of wavepackets $\phi_i^{A,\varphi}(t) = e^{-i\alpha \sin(\Omega t + \varphi)} \phi_i^B(t)$ with φ uniformly distributed in $[0, 2\pi)$, giving rise to the coherences

$$\frac{1}{2\pi} \int_0^{2\pi} \langle \phi_0^B, \phi_0^{A,\varphi} \rangle \langle \phi_A^{A,\varphi}, \phi_0^B \rangle \, d\varphi = \frac{1}{2\pi} \int_0^{2\pi} \int_0^{\infty} \int_0^{\infty} \Gamma^2 e^{-\Gamma(t+t')} e^{i\alpha(\sin(\Omega t + \varphi) - \sin(\Omega t' + \varphi))} \, dt' \, dt \, d\varphi$$

$$= \Gamma \int_0^{\infty} e^{-\Gamma t} J_0\left(2\alpha \sin\left(\frac{\Omega}{2}t\right)\right) dt \approx 1 - \frac{\alpha^2 \Omega^2}{2(\Gamma^2 + \Omega^2)} + O(\alpha^4).$$

(4.66)

From eq. (4.55), we thus obtain the Bell-state infidelity from periodic phase modulation, for small modulation indices α, as

$$\mathcal{E}_{\text{PM}} \approx \frac{\Omega^2}{\Gamma^2 + \Omega^2} \left(\frac{\alpha}{2}\right)^2 + O(\alpha^4).$$

(4.67)

The full expression, along with this approximation, is shown in fig. 4.23. Assuming that the modulation is caused by excess micromotion due to a stray electric field E_{stray} in the imaging direction, and the confinement with trap frequency ω_m in that direction to be given entirely by the RF pseudopotential, the infidelity due to micromotion is

$$\mathcal{E}_{\text{mm}} \approx \frac{1}{2} \frac{1}{\Gamma^2 + \Omega_{\text{rf}}^2} \left(\frac{ek_n}{m\omega_m} E_{\text{stray}} \right)^2 + \mathrm{O}\left(E_{\text{stray}}{}^4 \right), \tag{4.68}$$

where e is the ion charge, m the ion mass. For $^{88}\text{Sr}^+$ in a trap with $\Omega_{\text{rf}} = 2\pi \times 50\,\text{MHz}$ and a radial mode frequency of $\omega_m = 2\pi \times 3\,\text{MHz}$ aligned with the collection direction, an error of $\mathcal{E}_{\text{mm}} = 0.5\,\%$ is reached at a stray field of $E_{\text{stray}} = 39\,\text{Vm}^{-1}$. Here, the coupling efficiency has reduced less than 0.1 % (see fig. 4.19), confirming that any modulation of the probability envelope can safely be neglected.

In principle, the secular motion of the ion in the trap would similarly cause an error by reducing the wavefunction overlap, as described by eqs. (4.66) and (4.67). However, compared to the RF case, this will be suppressed by the lower motional frequency $\Omega = \omega_m$ (usually $\omega_m \ll \Gamma$). Laser cooling typically leaves the ion in close to a thermal state with some average motional occupation number $\bar{n}$. Considering the ions to follow a classical harmonic oscillator trajectory[20] after the reservoir interaction is "turned off" [Win+98], we can estimate the error by averaging over the decomposition into coherent states $|\tilde{\alpha}\rangle\langle\tilde{\alpha}|$ according to eq. (2.32). The dimensionless, complex coherent state amplitude $\tilde{\alpha}$ is related to the classical oscillation amplitude by $\langle\tilde{\alpha}|x|\tilde{\alpha}\rangle = 2\,x_0\,\text{re}(\tilde{\alpha})$, where x_0 is the spread of the ground-state wavefunction (see eq. (2.27)), giving the corresponding modulation index as $\alpha = 2k_n x_0 |\tilde{\alpha}|$. Hence, we obtain an estimate for the error due to thermal motion as

$$\mathcal{E}_{\bar{n}} = \frac{2}{\bar{n}} \int_0^\infty e^{-\frac{\tilde{\alpha}^2}{\bar{n}}} \mathcal{E}_{\text{PM}}\, \tilde{\alpha}\, d\tilde{\alpha} \approx \frac{\omega_m^2}{\Gamma^2 + \omega_m^2} \left(\frac{\alpha_{\bar{n}}}{2} \right)^2 + \mathrm{O}\left(\alpha_{\bar{n}}{}^4 \right), \tag{4.69}$$

where we have introduced an effective modulation index $\alpha_{\bar{n}} := 2k_n \sqrt{\bar{n}} x_0$. For the above case of $^{88}\text{Sr}^+$ with a radial mode frequency of $\omega_m = 2\pi \times 3\,\text{MHz}$, an error of $\mathcal{E}_{\bar{n}} = 0.5\,\%$ is reached at $\bar{n} \approx 62$ (while the Doppler cooling limit is $\bar{n} \approx 4$).

4.5.2   Residual birefringence

In the last section, we considered various situations with an ideal optical apparatus, where the ion–ion density matrices are invariably of the form (4.51). We now turn our attention to imperfections caused by the optical apparatus, for which this is no longer the case.

One instance of this is the misalignment of the two (still orthogonal) photon polarisations corresponding to the atomic σ^+ and π decay channels from one ion as compared

---

[20] In this estimate, we neglect the intrinsic micromotion which contributes an additional modulation at Ω_{rf}, see §2.2.1.

to the basis defined by the polarising beamsplitters. The stress birefringence in the optical fibres connecting Alice and Bob to the heralding station is not well-controlled, so a pair of wave plates at the input to the entangler is used to ensure the $\sigma \leftrightarrow H, \pi \leftrightarrow V$ polarisation mapping we assumed in eq. (4.46). If the calibration is erroneous (e.g. because the overall birefringence through the fibres slowly drifted over time), this potentially leads to an error in the entangled state. Assume Bob's polarisation states are rotated by a unitary U_{fibre}, while Alice's still follow the ideal mapping. The total system is then described by

$$U_P = U_P^{\text{ideal}} \cdot \mathcal{V}_{BH \leftrightarrow BV}(U_{\text{fibre}}) \tag{4.70}$$

with U_P^{ideal} from eq. (4.49).

The effect of birefringence that just adds a phase shift between the components, but leaves the $\sigma \leftrightarrow H, \pi \leftrightarrow V$ mapping unaffected, is trivial; it just shifts the entangled state phase. Explicitly, if $U_{\text{fibre}} = R_z(\theta) = \text{diag}(1, e^{i\theta})$, then

$$\rho_{\wedge p} = \frac{\eta^A \eta^B (\eta_{\text{det}})^2}{8} |\psi\rangle\langle\psi|, \quad |\psi\rangle = \frac{1}{\sqrt{2}} \left(|01\rangle \pm e^{-i\theta} |10\rangle \right), \tag{4.71}$$

with the plus sign for the same-side patterns $p \in \{A \wedge C, B \wedge D\}$, the minus sign for the opposite-side patterns $p \in \{A \wedge D, B \wedge C\}$. It can easily be corrected for by incorporating the phase shift into the later ion operations – indeed, in our experimental apparatus, we cannot even correct for this optically given that we only employ a pair of wave plates per input; rather, we recalibrate the ion–ion state phase assumed for subsequent operations as necessary.

As rotations around axes in the xy plane are equivalent up to conjugation by z rotations, we just consider rotations around the y axis, with

$$U_{\text{fibre}} = \begin{pmatrix} \cos \vartheta & \sin \vartheta \\ -\sin \vartheta & \cos \vartheta \end{pmatrix}. \tag{4.72}$$

As such a rotation mixes photons of different transitions, the results depend on the frequency splitting between the σ and π decays paths. Introducing the normalised frequency
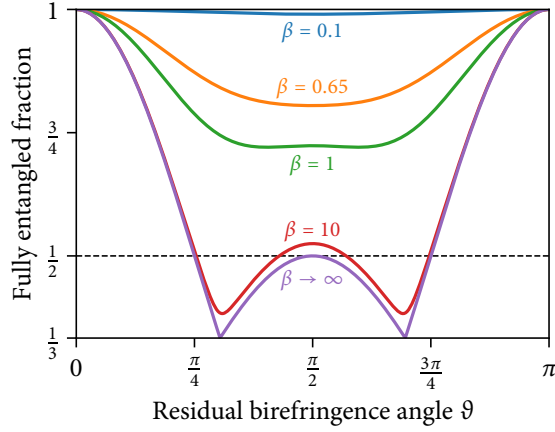
Figure 4.24: Loss in fidelity from uncompensated birefringence in the optical path connecting one of the nodes to the heralding station. Shown is the fully entangled fraction $\mathcal{F}$ for the ion–ion states post-selected on any of the four valid heralding patterns, as affected by a rotation with angle ϑ around any axis in the $x-y$ plane containing the circular and diagonal polarisation states. The strength of the effect depends on the normalised transition frequency splitting $\beta = \frac{(\nu_{e1}-\nu_{e0})B}{\Gamma}$; $\beta = 0.65$ matches the experiment.

splitting $\beta := \frac{(\nu_{e1}-\nu_{e0})B}{\Gamma}$, the time-averaged density matrices are

$$
\rho_{\wedge A \wedge C} = \rho_{\wedge B \wedge D} = \frac{\eta^A \eta^B (\eta_{\text{det}})^2}{16} \cdot
$$
$$
\begin{pmatrix}
(\sin\vartheta)^2 & -\frac{\sin\vartheta\cos\vartheta}{1-i\beta} & -\frac{\sin\vartheta\cos\vartheta}{1-i\beta} & -\frac{(\sin\vartheta)^2}{1-2i\beta} \\
-\frac{\sin\vartheta\cos\vartheta}{1+i\beta} & (\cos\vartheta)^2 & (\cos\vartheta)^2 & \frac{\sin\vartheta\cos\vartheta}{1-i\beta} \\
-\frac{\sin\vartheta\cos\vartheta}{1+i\beta} & (\cos\vartheta)^2 & (\cos\vartheta)^2 & \frac{\sin\vartheta\cos\vartheta}{1-i\beta} \\
-\frac{(\sin\vartheta)^2}{1+2i\beta} & \frac{\sin\vartheta\cos\vartheta}{1+i\beta} & \frac{\sin\vartheta\cos\vartheta}{1+i\beta} & (\sin\vartheta)^2
\end{pmatrix},
$$
$$
\rho_{\wedge A \wedge D} = \rho_{\wedge B \wedge C} = \frac{\eta^A \eta^B (\eta_{\text{det}})^2}{16} \cdot
$$

(4.73)

$$
\begin{pmatrix}
(\sin\vartheta)^2 & -\frac{\sin\vartheta\cos\vartheta}{1-i\beta} & \frac{\sin\vartheta\cos\vartheta}{1-i\beta} & \frac{(\sin\vartheta)^2}{1-2i\beta} \\
-\frac{\sin\vartheta\cos\vartheta}{1+i\beta} & (\cos\vartheta)^2 & -(\cos\vartheta)^2 & -\frac{\sin\vartheta\cos\vartheta}{1-i\beta} \\
\frac{\sin\vartheta\cos\vartheta}{1+i\beta} & -(\cos\vartheta)^2 & (\cos\vartheta)^2 & \frac{\sin\vartheta\cos\vartheta}{1-i\beta} \\
\frac{(\sin\vartheta)^2}{1+2i\beta} & -\frac{\sin\vartheta\cos\vartheta}{1+i\beta} & \frac{\sin\vartheta\cos\vartheta}{1+i\beta} & (\sin\vartheta)^2
\end{pmatrix}.
$$

The four patterns are all still equally likely to occur. If the frequency splitting is negligible compared to the transition linewidth, $|\beta| \ll 1$, then the result is just a maximally entangled state in a rotated basis,

$$
\rho_{\wedge P}^{\beta=0} = \frac{\eta^A \eta^B (\eta_{\text{det}})^2}{8} |\psi\rangle\langle\psi|, \quad |\psi\rangle := \frac{1}{\sqrt{2}} \left(-\sin\vartheta |00\rangle + \cos\vartheta |01\rangle \pm \cos\vartheta |10\rangle \pm \sin\vartheta |11\rangle \right),
$$

(4.74)

where the plus signs hold apply for the same-side patterns $p \in \{A \wedge C, B \wedge D\}$, the minus signs for the opposite-side patterns $p \in \{A \wedge D, B \wedge C\}$.

If the frequency splitting cannot be neglected, the states are no longer maximally entangled (in our apparatus, $B \approx 0.5\,\text{mT}$, so in $^{88}\text{Sr}^+$, $\beta \approx 0.65$). The fully entangled fraction

for any of the patterns, as a function of the rotation angle ϑ for any axis in the $x - y$ plane, is shown in fig. 4.24. To yield an infidelity of $\mathcal{E}^{\vartheta,\beta=0.65} < 0.5\,\%$, the amount of rotation must be less than $|\vartheta| < 0.13 = 7°$.[21]

4.5.3 Imperfect 50 : 50 beamsplitter

Next, we consider a deviation of the non-polarising beamsplitter from the ideal $50 : 50$ splitting ratio, as expressed by a deviation of the angles in eq. (4.49) from $\theta_H^{50:50} = \theta_V^{50:50} = \frac{\pi}{4}$. Here, the consequences qualitatively depend on how the error affects the two polarisations (there are no temporal effects, so we directly consider the detection-time-averaged states):

Let $\theta_H^{50:50} = \frac{\pi}{4} + \delta_H$, $\theta_V^{50:50} = \frac{\pi}{4} + \delta_V$. The sub-normalised density matrices are then

$$\rho_{\wedge A \wedge C} = \frac{\eta^A \eta^B (\eta_{\det})^2}{16} \cdot$$
$$\begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & (1 - \sin(2\delta_H))(1 + \sin(2\delta_V)) & \frac{\cos(2(\delta_H+\delta_V))}{2} + \frac{\cos(2(\delta_H-\delta_V))}{2} & 0 \\ 0 & \frac{\cos(2(\delta_H+\delta_V))}{2} + \frac{\cos(2(\delta_H-\delta_V))}{2} & (1 + \sin(2\delta_H))(1 - \sin(2\delta_V)) & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix},$$

$$\rho_{\wedge B \wedge D} = \frac{\eta^A \eta^B (\eta_{\det})^2}{16} \cdot$$
$$\begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & (1 + \sin(2\delta_H))(1 - \sin(2\delta_V)) & \frac{\cos(2(\delta_H+\delta_V))}{2} + \frac{\cos(2(\delta_H-\delta_V))}{2} & 0 \\ 0 & \frac{\cos(2(\delta_H+\delta_V))}{2} + \frac{\cos(2(\delta_H-\delta_V))}{2} & (1 - \sin(2\delta_H))(1 + \sin(2\delta_V)) & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix},$$

$$\rho_{\wedge A \wedge D} = \frac{\eta^A \eta^B (\eta_{\det})^2}{16} \cdot$$
$$\begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & (1 - \sin(2\delta_H))(1 - \sin(2\delta_V)) & -\frac{\cos(2(\delta_H+\delta_V))}{2} - \frac{\cos(2(\delta_H-\delta_V))}{2} & 0 \\ 0 & -\frac{\cos(2(\delta_H+\delta_V))}{2} - \frac{\cos(2(\delta_H-\delta_V))}{2} & (1 + \sin(2\delta_H))(1 + \sin(2\delta_V)) & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix},$$

$$\rho_{\wedge B \wedge C} = \frac{\eta^A \eta^B (\eta_{\det})^2}{16} \cdot$$
$$\begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & (1 + \sin(2\delta_H))(1 + \sin(2\delta_V)) & -\frac{\cos(2(\delta_H+\delta_V))}{2} - \frac{\cos(2(\delta_H-\delta_V))}{2} & 0 \\ 0 & -\frac{\cos(2(\delta_H+\delta_V))}{2} - \frac{\cos(2(\delta_H-\delta_V))}{2} & (1 - \sin(2\delta_H))(1 - \sin(2\delta_V)) & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}.$$

$$(4.75)$$

---

[21] As a reminder, this only captures the decoherence resulting from the frequency mismatch; any form of birefringence will also change the form of the maximally entangled state produced, cf. eqs. (4.71) and (4.74)).

If the ratio differs from ideal the same for both polarisations, $\delta_H = \delta_V = \delta$, then the heralding patterns with clicks on the same side of the beamsplitter, stay unaffected except for a slight decrease in probability

$$\mathcal{F}(\rho_{|(A\wedge C)}^{\delta_H=\delta_V}) = \mathcal{F}(\rho_{|(B\wedge D)}^{\delta_H=\delta_V}) = 1,$$
$$\operatorname{tr}\rho_{\wedge A\wedge C}^{\delta_H=\delta_V} = \operatorname{tr}\rho_{\wedge B\wedge D}^{\delta_H=\delta_V} = \frac{\eta^A\eta^B(\eta_{\det})^2}{8}\cos(2\delta)^2, \tag{4.76}$$

but for the patterns on opposite sides of the beamsplitter, the fidelity is lowered but the probability is increased:

$$\mathcal{F}(\rho_{|(A\wedge D)}^{\delta_H=\delta_V}) = \mathcal{F}(\rho_{|(B\wedge C)}^{\delta_H=\delta_V}) = \frac{1}{1+\sin(2\delta)^2} \approx 1 - (2\delta)^2 + O(\delta^3),$$
$$\operatorname{tr}\rho_{\wedge A\wedge D}^{\delta_H=\delta_V} = \operatorname{tr}\rho_{\wedge B\wedge C}^{\delta_H=\delta_V} = \frac{\eta^A\eta^B(\eta_{\det})^2}{8}\frac{1+\sin(2\delta)^2}{2}. \tag{4.77}$$

The situation is reversed if the error is antisymmetric in the polarisations, $\delta_H = \delta, \delta_V = -\delta$. Now, the opposite-side patterns are unaffected except for a decrease in probability,

$$\mathcal{F}(\rho_{|(A\wedge D)}^{\delta_H=-\delta_V}) = \mathcal{F}(\rho_{|(B\wedge C)}^{\delta_H=-\delta_V}) = 1,$$
$$\operatorname{tr}\rho_{\wedge A\wedge D}^{\delta_H=-\delta_V} = \operatorname{tr}\rho_{\wedge B\wedge C}^{\delta_H=-\delta_V} = \frac{\eta^A\eta^B(\eta_{\det})^2}{8}\cos(2\delta)^2 \tag{4.78}$$

while the same-side patterns are affected,

$$\mathcal{F}(\rho_{|(A\wedge D)}^{\delta_H=-\delta_V}) = \mathcal{F}(\rho_{|(B\wedge C)}^{\delta_H=-\delta_V}) = \frac{1}{1+\sin(2\delta)^2} \approx 1 - (2\delta)^2 + O(\delta^3),$$
$$\operatorname{tr}\rho_{\wedge A\wedge C}^{\delta_H=-\delta_V} = \operatorname{tr}\rho_{\wedge B\wedge D}^{\delta_H=-\delta_V} = \frac{\eta^A\eta^B(\eta_{\det})^2}{8}\frac{1+\sin(2\delta)^2}{2}. \tag{4.79}$$

In both these cases, the overall coincidence probability remains the same, but the fidelity averaged across all patterns (weighted by the respective probability) drops by

$$\mathcal{E}_\delta^{\delta_H=\pm\delta_V} = \frac{\sin(2\delta)^2}{2}, \tag{4.80}$$

though this could of course be avoided at the cost of a factor > 2 in success probability by not using the affected click patterns. Both Hong–Ou–Mandel-suppressed patterns occur with the same probability in either case,

$$\rho_{\wedge A\wedge B}^{\delta_H=\pm\delta_V} = \frac{\eta^A\eta^B(\eta_{\det})^2}{4}\sin(2\delta)^2\,|0\rangle\langle 0|,$$
$$\rho_{\wedge C\wedge D}^{\delta_H=\pm\delta_V} = \frac{\eta^A\eta^B(\eta_{\det})^2}{4}\sin(2\delta)^2\,|1\rangle\langle 1|. \tag{4.81}$$

For $\mathcal{E}_\delta^{\delta_H=\pm\delta_V} < 1\,\%$, the imbalance must be less than $|\delta| < 0.071$, equivalent to a $57:43$ beamsplitter ratio.

If the imbalance does not affect the polarisations (anti-)symmetrically, none of the patterns remain unaffected in fidelity. For instance, if $\delta_V = 0$, then

$$\mathcal{F}(\rho_{|(A \wedge C)}^{\delta_V=0}) = \mathcal{F}(\rho_{|(B \wedge D)}^{\delta_V=0}) = \mathcal{F}(\rho_{|(A \wedge D)}^{\delta_V=0}) = \mathcal{F}(\rho_{|(B \wedge C)}^{\delta_V=0}) = \frac{1}{2}\left(1 + |\cos(2\delta_H)|\right),$$

$$\operatorname{tr} \rho_{\wedge A \wedge C}^{\delta_V=0} = \operatorname{tr} \rho_{\wedge B \wedge D}^{\delta_V=0} = \operatorname{tr} \rho_{\wedge A \wedge D}^{\delta_V=0} = \operatorname{tr} \rho_{\wedge B \wedge C}^{\delta_V=0} = \frac{\eta^A \eta^B (\eta_{\det})^2}{8}.$$

$$(4.82)$$

In this case, only the Hong–Ou–Mandel-suppressed coincidences corresponding to the affected polarisation (here H) is observed, with frequency still given by eq. (4.81).

4.5.4 Imperfect polarising beamsplitters

Finally, we consider issues caused by the polarising beamsplitters.

PBS extinction ratio. We model polarising beamsplitters that do not perfectly separate the polarisations, but leave a certain amount of V polarisation in the transmitted beam or H polarisation in the reflected beam, by setting the angles[22] $\theta_H^{\text{PBS A}} = \delta_H^A$, $\theta_H^{\text{PBS B}} = \delta_H^B$, $\theta_V^{\text{PBS A}} \frac{\pi}{2} + \delta_V^A$, $\theta_V^{\text{PBS B}} \frac{\pi}{2} + \delta_V^B$ in eq. (4.49). The effects do not depend on the temporal wavepackets, so we directly consider the detection-time-averaged states:

$$\rho_{\wedge A \wedge C} = \frac{\eta^A \eta^B (\eta_{\det})^2}{32}.$$

$$\begin{pmatrix} 1 - \cos(4\delta_H^A) & 0 & 0 & 0 \\ 0 & 1 + \cos(2\delta_H^A)\cos(2\delta_V^A) & 1 + \cos(2\delta_H^A)\cos(2\delta_V^A) & 0 \\ 0 & 1 + \cos(2\delta_H^A)\cos(2\delta_V^A) & 1 + \cos(2\delta_H^A)\cos(2\delta_V^A) & 0 \\ 0 & 0 & 0 & 1 - \cos(4\delta_V^A) \end{pmatrix},$$

$$\rho_{\wedge A \wedge C} = \frac{\eta^A \eta^B (\eta_{\det})^2}{32}.$$

$$\begin{pmatrix} 1 - \cos(4\delta_H^B) & 0 & 0 & 0 \\ 0 & 1 + \cos(2\delta_H^B)\cos(2\delta_V^B) & 1 + \cos(2\delta_H^B)\cos(2\delta_V^B) & 0 \\ 0 & 1 + \cos(2\delta_H^B)\cos(2\delta_V^B) & 1 + \cos(2\delta_H^B)\cos(2\delta_V^B) & 0 \\ 0 & 0 & 0 & 1 - \cos(4\delta_V^B) \end{pmatrix},$$

$$\rho_{\wedge A \wedge D} = \frac{\eta^A \eta^B (\eta_{\det})^2}{8}\left((\cos\delta_H^A)^2(\cos\delta_V^B)^2 + (\sin\delta_V^A)^2(\sin\delta_H^B)^2\right)|\Psi^-\rangle\langle\Psi^-|,$$

$$\rho_{\wedge B \wedge C} = \frac{\eta^A \eta^B (\eta_{\det})^2}{8}\left((\cos\delta_V^A)^2(\cos\delta_H^B)^2 + (\sin\delta_H^A)^2(\sin\delta_V^B)^2\right)|\Psi^-\rangle\langle\Psi^-|.$$

$$(4.83)$$

The detection patterns with detectors on opposite sides of the 50 : 50 beamsplitter are

---

[22]Particularly in the regime of very high extinction ratios, the remaining imperfections might well be non-uniform over the beam diameter, which would result in a non-unitary transformation once the spatial degree is traced out. As we do not further modify the polarisations, however, but merely direct the outputs onto photodetectors, that distinction is without consequence here.
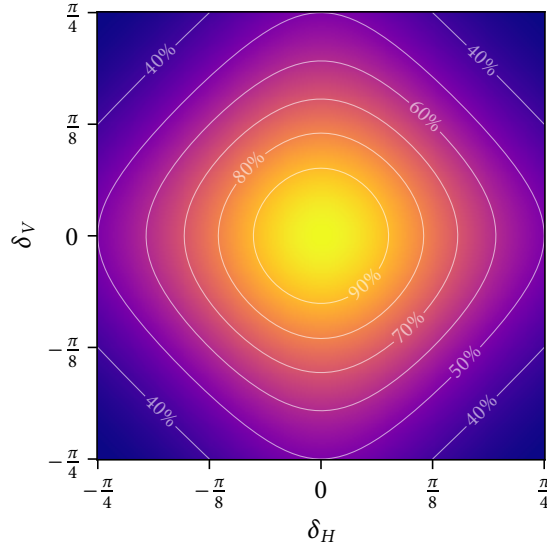
Figure 4.25: Fully entangled fraction of the ion–ion state post-selected on the detection of two photons on the outputs of a polarising beamsplitter with imperfect extinction for either polarisation, as given by the angles δ_H and δ_V specifying the deviation from perfectly transmitting H polarisation and perfectly reflecting V polarisation.

unaffected by any PBS imperfections except for a reduction in rate; the Hong–Ou–Mandel effect alone is enough to uniquely project the photonic part into the antisymmetric Bell state. The same-side patterns increase in rate compensating the decrease in the others, but drop in fidelity for an imperfect PBS on the respective side. Dropping the superscripts identifying the side,

$$\mathcal{E}_{A\wedge C, B\wedge D}^{\delta_H, \delta_V} = \frac{1}{1 - 2\frac{\cos(2\delta_H)\cos(2\delta_V)+1}{\cos(4\delta_H)+\cos(4\delta_V)-2}} \approx 2\,\delta_H{}^2 + 2\,\delta_V{}^2 + O(\delta_H{}^4 + \delta_V{}^4), \tag{4.84}$$

which is depicted in fig. 4.25. The error in PBSes is often asymmetric, but if we assume $\delta_H^s = \delta_V^s = \delta$ for a first estimate, the average over all patterns (weighted by the respective probability) is

$$\mathcal{E}_\delta = \frac{\sin(2\delta)^2}{2}. \tag{4.85}$$

For the infidelity to remain below $\mathcal{E}_\delta < 1\,\%$, $|\delta| < 0.071$, i.e. the PBSes are required to have an extinction ratio better than $200 : 1$.

Tilted PBS axes. Even for a perfectly polarising beamsplitter, a last source of imperfections lies in the orientation of the polarisation measurement bases. If the basis set by one PBS is used to calibrate out the fibre birefringence, and the measurement axis set by the other PBS is rotated with respect to it – whether due to it being mounted at an angle, due to manufacturing differences, or different birefringence in the two paths following the $50 : 50$ beamsplitter – heralding patterns involving that second PBS will be negatively impacted.

Concretely, assume that PBS B is rotated by an angle α around the optical axis compared to the other PBS and the ion decay channel mapping. Retaining the indices H and V for the transmitted and reflected polarisations of the beamsplitter for ease of notation, the total unitary is

$$
\begin{aligned}
U_P = \ & \mathcal{V}_{BV \leftrightarrow DV} \left(U_{\mathrm{BS}}(\theta_V^{\mathrm{PBS\ B}}) \right) \cdot \mathcal{V}_{BH \leftrightarrow DH} \left(U_{\mathrm{BS}}(\theta_H^{\mathrm{PBS\ B}}) \right) \cdot \\
& \mathcal{V}_{BH \leftrightarrow BV} \left(U_{\mathrm{rot}}(\alpha) \right) \cdot \\
& \mathcal{V}_{AV \leftrightarrow CV} \left(U_{\mathrm{BS}}(\theta_V^{\mathrm{PBS\ A}}) \right) \cdot \mathcal{V}_{AH \leftrightarrow CH} \left(U_{\mathrm{BS}}(\theta_H^{\mathrm{PBS\ A}}) \right) \cdot \\
& \mathcal{V}_{AV \leftrightarrow BV} \left(U_{\mathrm{BS}}(\theta_V^{50:50}) \right) \cdot \mathcal{V}_{AH \leftrightarrow BH} \left(U_{\mathrm{BS}}(\theta_H^{50:50}) \right)
\end{aligned}
\tag{4.86}
$$

with

$$
U_{\mathrm{rot}}(\alpha) = \begin{pmatrix} \cos\alpha & \sin\alpha \\ -\sin\alpha & \cos\alpha \end{pmatrix}.
\tag{4.87}
$$

The same-side pattern on the other PBS ($A \wedge C$) is entirely unaffected. The two opposite-side patterns still lead to the same, maximally-entangled post-selected states, as guaranteed by the Hong–Ou–Mandel effect. Their probability drops, however:

$$
\operatorname{tr} \rho_{\wedge A \wedge D} = \operatorname{tr} \rho_{\wedge B \wedge C} = \frac{\eta^A \eta^B (\eta_{\mathrm{det}})^2}{8} (\cos\alpha)^2.
\tag{4.88}
$$

Lastly, the phases in the ion–ion state post-selected on two clicks on the ports of the rotated PBS become dependent on the detection time (unless α is a multiple of $\pi/2$), such that the averaged density matrix is

$$
\rho_{\wedge B \wedge D}^{\alpha} = \frac{\eta^A \eta^B (\eta_{\mathrm{det}})^2}{32} \cdot
\begin{pmatrix}
1 - \cos(4\alpha) & -\frac{\sin(4\alpha)}{1 - i\beta} & -\frac{\sin(4\alpha)}{1 - i\beta} & \frac{\cos(4\alpha) - 1}{1 - 2i\beta} \\
-\frac{\sin(4\alpha)}{1 + i\beta} & 1 + \cos(4\alpha) & 1 + \cos(4\alpha) & \frac{\sin(4\alpha)}{1 - i\beta} \\
-\frac{\sin(4\alpha)}{1 + i\beta} & 1 + \cos(4\alpha) & 1 + \cos(4\alpha) & \frac{\sin(4\alpha)}{1 - i\beta} \\
\frac{\cos(4\alpha) - 1}{1 + 2i\beta} & \frac{\sin(4\alpha)}{1 + i\beta} & \frac{\sin(4\alpha)}{1 + i\beta} & 1 - \cos(4\alpha)
\end{pmatrix},
\tag{4.89}
$$

where $\beta := \frac{(\nu_{e1} - \nu_{e0})B}{\Gamma}$ is again the Zeeman frequency splitting normalised by the excited state linewidth. In fact, $\rho_{\wedge B \wedge D}^{\alpha}$ is very similar in form to the result for uncompensated birefringence from eq. (4.73), with $\vartheta = 2\alpha$. The fully entangled fraction behaves exactly as illustrated for that case in fig. 4.24: if α is a multiple of $\pi/2$, then the atomic decay channels map exactly to the PBS axes (possibly swapped from the intended mapping), and no reduction in fidelity occurs. For the one pattern to lose less than 1 % in fidelity, $|\alpha| < 5°$ (or $|\alpha| < 11°$ for less than 1 % loss when averaged over all four patterns).

4.5.5 Imperfections not considered

In the above, we have considered the single-photon detectors to be perfect (if of limited quantum efficiency). In particular, we have neglected dark counts, as they are negligible in

our experiment; across a typical 30 ns heralding window, the dark count rate of 50 Hz of our detectors (which is fairly high compared to the best-in-class performance) only yields a 2×10^{-6} probability for a false click, compared to 2×10^{-2} to detect a photon from an ion. We have also neglected the timing resolution and jitter of the detectors – while it leads to dephasing in ion–photon experiments, the ion–ion state is, in the ideal configuration, time-independent, and we do not rely on tight windowing to combat arrival time mismatch or dark counts. If of interest, both of these effects would be readily integrated into the model by replacing the detection operator by the more complex POVM elements from ref. [Gou+18], however[23].

We also have not considered any non-unitary effects, such as due to imperfect overlap of the input modes on the non-polarising beamsplitter, non-uniformities (e.g. in birefringence) across the beam modes, or differential loss between the polarisation components. Such effects could be described by generalising the conjugation of $\rho^A_{I^A P} \otimes \rho^B_{I^B P}$ by $\mathbb{1}_{I^A I^B} \otimes U_P$ in eq. (4.48) to the application of a general completely positive map.

4.6   Summary

We have studied the generation of ion–photon entanglement through spontaneous decay of a trapped ion after excitation to a short lived state with σ^+ and π dipole-allowed decay channels, for instance on the $S_{1/2} \leftrightarrow P_{1/2}$ transition in $^{88}\mathrm{Sr}^+$ (and other commonly used species such as $^{40}\mathrm{Ca}^+$, $^{138}\mathrm{Ba}^+$).

The use of single-mode fibres as spatial mode filters was shown to select only contributions from the transversal components of the atomic dipole moments. This results in polarisation only parallel to that of the respective atomic dipole moment component at the fibre output. Conditioned on the detection of a photon, decays on the $S_{1/2} \leftrightarrow P_{1/2}$ transition thus give rise to maximally entangled ion–photon states.

This result is valid for arbitrarily large collection angles, which allows the implementation of a coherent ion–photon interface with high heralded success probability, and, owing to the simple excitation scheme in $^{88}\mathrm{Sr}^+$, a relatively high attempt rate, using polarisation encoding. The latter is convenient for basic experiments, as polarisation states are readily manipulated using birefringent optics. In the half-space limit of NA 1.0, up to 22 % of the total emission can be coupled into a Gaussian single-mode fibre using a standard imaging system. We also derived an apparent shift in the emitter position for emission from transitions with non-zero angular momentum, and showed that the correctly aligned configuration can still be identified in the experiment by optimising the fibre position based on π emission only (which does not exhibit this shift), or on average unpolarised steady-state fluorescence.

---

[23] If dark counts are significant, it will also be necessary to include the $\rho^s_{\mathrm{dark},I^s}$ terms from eq. (4.45) for an accurate estimate of the post-herald state.

We also presented a numerical study of the system performance under small misalignments and other relevant experimental imperfections, and showed that fidelities well above 99.5 % should be reachable in a typical $^{88}Sr^+$ experiment with NA 0.6 collection optics. This is in contrast to a free-space fidelity limit of 93.5 % at the same numerical aperture due to polarisation mixing effects.

Finally, we described a generic method for treating the interference and detection of pairs of single photons in a linear-optics instrument such as the partial Bell-basis measurement setup used for heralded entanglement swapping, integrating entanglement with other subsystems (the ions) into the full density-operator treatment. We analytically calculated the post-herald ion–ion states and loss of Hong–Ou–Mandel suppression under a number of imperfections, including that of periodic phase modulation due to ion motion.

Most of the above considerations are not specific to the trapped-ion systems studied here, but also apply to quantum dipole emitters more generally, or in case of the heralded entanglement swapping procedure, to generic qubit–photon systems. As high quantum gate fidelities are also being reached in such platforms, such as in nitrogen-vacancy colour centres in diamond lattices, the results presented here should find wider applicability to networking with other qubit platforms as well.

5 Micromotion minimisation by synchronous detection of parametrically excited motion

An inherent feature of Paul traps, where a quadrupole potential oscillating in the radio-frequency range (RF) provides effective time-averaged confinement of charged particles [Pau90], is "micromotion": the coherent modulation of particle trajectories on the RF time scale, which is minimal[1] at the central null of the quadrupole potential. If the equilibrium position of the particle does not coincide with this null, the particle undergoes excess micromotion that cannot be reduced by cooling.

Excess micromotion can be avoided by applying electrostatic compensation fields that shift the equilibrium position onto the RF null, but for this, a way to detect this mismatch is necessary. Across the history of RF traps, a wide range of approaches have been proposed for this. However, as detailed in §5.1, none of them are particularly well suited to our case of a planar surface electrode trap where all the laser beams propagate parallel to the trap surface. To be able to compensate stray fields normal to the trap surface, we employ a novel method based on amplitude modulation of the trapping potential, which we introduce, along with a comprehensive theoretical treatment, in §5.2. We present experimental data validating the theoretical model (§5.3), and discuss some practical concerns regarding the use of this method to build a robust automated procedure for the compensation of stray fields in typical non-ideal surface traps without detailed knowledge about the motional mode structure (§5.4). We discuss the statistical floor to the practically achievable compensation accuracies in our setup (§5.5), and compare the result to two other micromotion detection methods, for which we observe significant systematic shifts (§5.6). Our method is robust, achieves a high sensitivity, and requires a minimal degree of technical complexity. This makes it well-suited for many other applications of Paul traps in quantum information processing and beyond; we conclude the chapter with a discussion of its merits compared to other techniques.

---

[1] Even if the average position of the particle does coincide with the RF null, the trajectories of secular motion with non-zero amplitude will still be modulated at the trap drive frequency in what is known as intrinsic micromotion; in a quantum description, the harmonic oscillator wavepackets will "breathe" accordingly even in the ground state (see §2.2.1).

5.1   Background

Uncontrolled excess micromotion is rarely a desirable feature. For atomic clocks using trapped ions, the modulation in velocity leads to an appreciable time-dilation shift in the observed frequency; the associated uncertainty has dominated the accuracy budget of recent experiments [Bre+19]. In quantum information experiments, micromotion can not only lead to increased sensitivity of the heating rate of the secular modes of motion to RF noise [BP01], but also introduces a periodic phase modulation to all interactions of the trapped ions with external fields (such as lasers or local microwave gradients). In our case, we are also interested in minimising micromotion in the direction normal to the trap surface ("out of plane", OOP) to avoid phase modulation of the photons entangled with the ion state.

Trap electrode geometries are usually designed such that the RF and DC fields both vanish at the intended trapping location[2]. For a single trapped particle, excess micromotion is thus the consequence of imperfections in the experimental realisation, such as DC field offsets caused by stray charges on nearby elements of the trap apparatus, or phase shifts between different RF electrodes. Here, we focus on the former, which is generally the larger source of micromotion. It can be counteracted by shifting the DC field null back onto the RF null using static compensation potentials applied to a suitable combination of trap electrodes.

Various techniques to determine the required compensation fields have been proposed. Some of these micromotion compensation methods make use of particularities of the experimental systems for which they were developed, for instance a co-trapped gas of ultra-cold atoms [Här+13; Mei+16], a high-finesse optical cavity collecting ion emission [Chu+13], microwave near-field magnetic gradients [War+13], strongly-focussed high-power laser beams for optical dipole trapping [Hub+14], or fluorescence imaging optics with adjustable focus [Glo+15]. If the RF confinement strength can be varied, stray fields lead to a change in the ion equilibrium position, which can be detected using camera imaging [Ber02; Hir+20], or, as very recently demonstrated, with exquisite precision through coherent detection of the corresponding phase shift in a laser beam addressing a narrow-linewidth transition [Hig+21].

Another, generally applicable family of methods relies on micromotion-induced phase modulation of a laser beam in the ion rest frame [Ber+98]. These include minimisation of the amplitude of the micromotion modulation sideband observed for a transition with a linewidth smaller than the trap drive frequency, or minimising temporal variations in the stimulated emission rate for a transition of linewidth comparable to or larger than the

---

[2]In many linear trap designs, the finite extent of the trap leads to a small amount of axial micromotion everywhere except for one point of symmetry. Thus, the fields only completely vanish for a single confined ion; for strings of more than one ion, the ions away from the centre will experience a small but non-zero RF field even if the string is positioned exactly at the intended location.
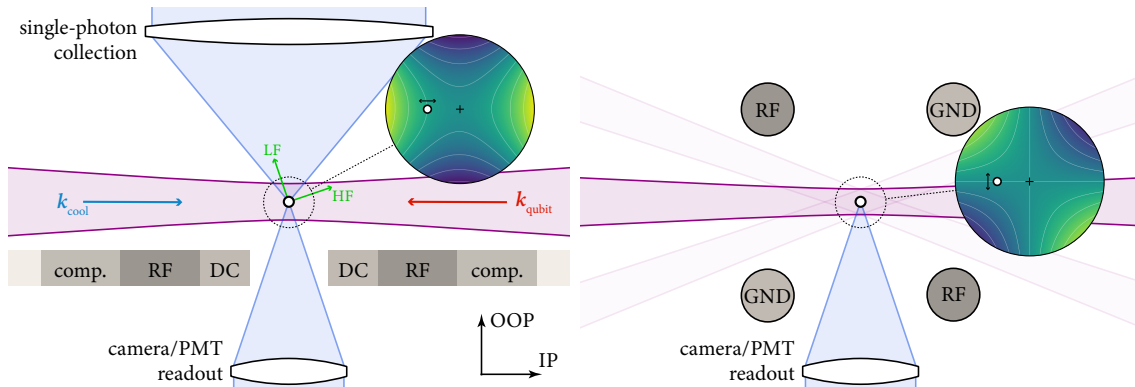
Figure 5.1: Radial slices through the *HOA2* surface trap used here (*left*), and a prototypical "four-rod" macroscopic Paul trap (*right*); not to scale. The ion position in the centre is illustrated by the small white dot. In the former case, all the laser beams are confined to the same plane parallel to the trap surface (including incoherent cooling beams $\boldsymbol{k}_{\text{cool}}$, and those for coherent qubit manipulation $\boldsymbol{k}_{\text{qubit}}$). The insets show a snapshot of the RF potential near the ion position, illustrating the different orientation of the principal axes, which leads to opposite relations of ion displacement vs. induced micromotion. For the surface trap, the radial motional modes of the ion in the pseudopotential are also shown (labelled LF and HF, referring to the lower- and higher-frequency one), which are the mechanism by which our method achieves directional resolution.

trap RF (see ref. [Kel+15] for a comprehensive discussion). Such methods measure the projection of micromotion onto the wavevector $\boldsymbol{k}$ of a probing laser. As such, multiple probe beams are required to be able to resolve stray fields in all directions.

5.1.1 *Limitations of previous methods in surface traps*

The methods based on micromotion-induced modulation of a laser beam are widely used throughout ion trap research. However, their intrinsic coupling between the probed direction and that of a laser beam is challenging in the planar "surface trap" designs produced through scalable microfabrication techniques, such as that used here (see chapter 3). There, it is usually not possible to achieve a direction $\boldsymbol{k}$ with an appreciable component perpendicular to the trap plane without the laser beam also striking the trap, which is generally undesirable due to charging effects (with the exception of the infrared beams necessary to cool some ion species [All+10]).

Our particular case is illustrated in fig. 5.1: The axis normal to the trap plane is taken up by the high-NA single-photon collection optics, and while there is access through the central slot below the ion position, it is taken up by the state-dependent fluorescence readout optics. All the laser beams propagate, to good approximation, exactly in the same plane parallel to the trap surface; consequently, methods based on laser phase modulation are not sensitive to micromotion in the out-of-plane direction. Furthermore, the principal axes of the RF potential minimum are rotated by 45° compared to a typical situation in a macroscopic Paul trap, such that a positional shift along the in-plane or out-of-plane directions (where there is camera and laser access) results in micromotion parallel to the

displacement, rather than normal to it. Consequently, a number of combinations which allow the probing of stray fields without extra laser access in four-rod-like geometries are not applicable here[3].

The key idea to address the challenge of stray field compensation in the out-of-plane direction without introducing any extra laser beams along that axis – which we would rather avoid, as the alignment requirements to achieve a high single-photon collection efficiency are already challenging without needing to take extra elements for additional beams into account – is to address it using a mechanism that already has a projection in the vertical direction: the radial modes of motion of the ion in the RF pseudopotential. These can be addressed in a way that is sensitive to the ion position w.r.t. the RF null through modulation of the RF trapping potential, as described in the next section.

5.2   *Parametric excitation*

Parametric excitation refers to the phenomenon that motion in a harmonic oscillator can be excited through periodic modulation of its parameters over time. In particular, if the oscillation frequency ω_0 is modulated at a frequency ω_{am} twice that of the unperturbed oscillator or integer subharmonics thereof ($\omega_{am} = 2\,\omega_0/n$ for $n \in \mathbb{N}$), a parametric resonance is triggered, where the oscillation amplitude grows exponentially over time (unless counteracted by a damping term).

Parametric instabilities of this kind have been studied in a wide variety of mechanical and electrical systems [RS16]. For ions confined in Paul traps in particular, parametric excitation has been used for the mass-selective removal of ion species from the trap [Sch+20], and also for stray field compensation [ITU11; Nar+11; Tan+12]. These proposals are based on the fact that the quickly-growing solutions of a system driven at parametric resonance can easily lead to large ion orbits where changes to the average fluorescence scattering rate during laser cooling can be observed. As the unstable behaviour in the presence of damping from laser cooling empirically somewhat depends on the distance from the ion equilibrium position to the RF null – that is, the strength of the stray field –, this can be used for micromotion minimisation by trying to determine the strength of applied compensation field for which the excitation vanishes. For that purpose, excitation at the first parametric resonance ($n = 1$, i.e. $\omega_{am} = 2\,\omega_0$) is most convenient, as the instability region is broadest there.

---

[3]For instance, in traps resembling a four-rod design, one could probe vertical fields with a horizontal laser beam, and horizontal fields by determining the position shift for different RF confinement strengths using a vertically-placed camera [Ber02], or even the same horizontal beam(s) using coherent techniques [Hig+21]. In our case, camera imaging covers the two directions already accessible via laser beams, while offering barely any resolution in the out-of-plane direction, at least without extra mechanical elements to scan a narrow plane of focus [Glo+15].
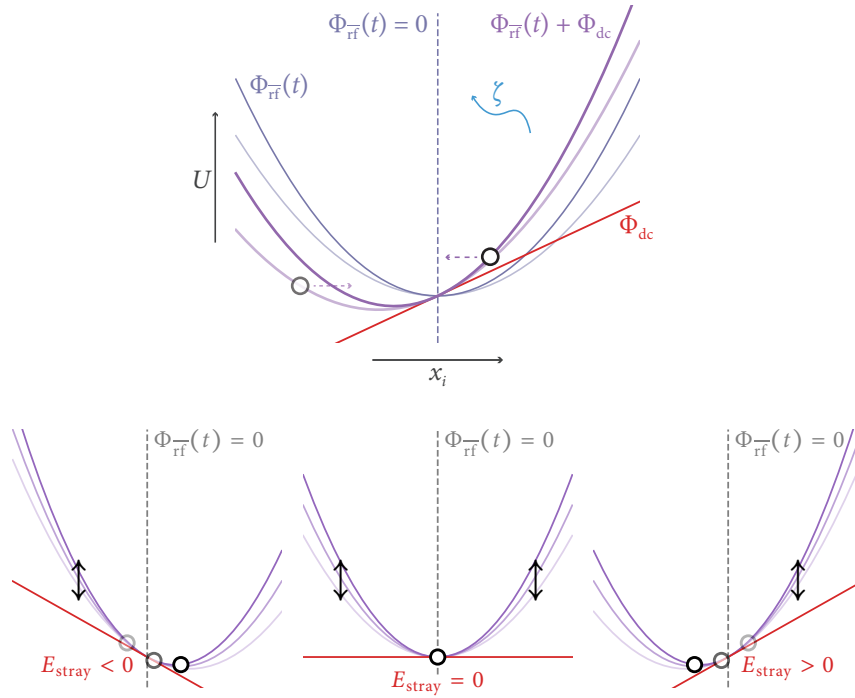
Figure 5.2: Principle of the parametric excitation method for stray field detection. *Top:* The total effective potential U experienced by the ion along one radial direction x_i is a sum of the averaged confining effect of the RF potential, expressed by the pseudopotential $\Phi_{\overline{rf}}$, and any static potential Φ_{dc}. Amplitude modulation of the RF potential near the frequency of the ion's secular mode of motion excites a parametric resonance, which in the presence of some linear damping ζ leads to steady-state oscillations. *Bottom:* The steady-state amplitude of the excited motion depends on the signed amplitude of the stray electric field E_{stray}.

These previous proposals already share an important characteristic with the method described here, in that they allow one to distinguish fields in multiple spatial directions by addressing the modes of motion in frequency space. They also have some practical shortcomings, however. As the amplitude of the observed signal is a result of various non-linear effects, it is not always easy to interpret; for instance, based on a number of parameters, the resonance can lead to an increase [ITU11] in observed count rate, but also a decrease [Nar+11]. Furthermore, as the signal results from what is fundamentally an unstable runaway process (limited only by non-idealities in the system), it is relatively easy to accidentally eject the ion from the trap, especially when working with large excitation amplitudes and little damping in order to obtain good resolution.

Here, we propose[4] to work in a completely different regime instead: near the second parametric resonance ($\omega_{am} \approx \omega_0$), but for a small modulation depth h such that damping from laser cooling suppresses the exponential growth in amplitude. As illustrated in

---

[4] The synchronous detection of motion excited through sidebands on the trap RF drive was already described in A. M. Eltony's 2013 MSc thesis as "synchronous tickle measurement" [Elt13]. However, the theoretical treatment there was limited, modelling the interaction as an oscillator driven by a uniform force rather than a parametric interaction (ignoring the potential of the latter for instability) among a number of other simplifications. To my knowledge, these earlier efforts were – unfortunately! – never formally published.
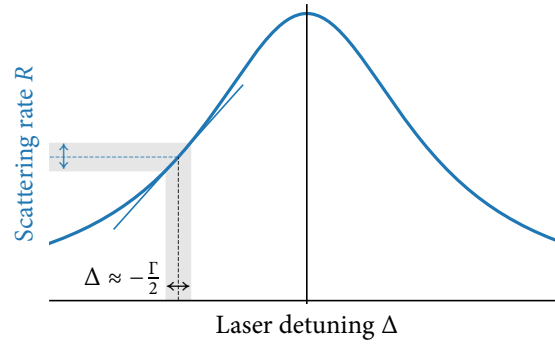
Figure 5.3: The parametrically excited motion can be detected through modulation of the ion fluorescence rate if the laser parameters are chosen such that there is a linear dependence of scattering rate to the laser detuning (e.g. at a detuning of half the transition linewidth Γ). Then, the Doppler effect causes parametrically driven oscillations at ω_{am} to produce a modulation in scattering rate at the same frequency ω_{am}. Thus, as long as the wavevector of the cooling beam has a non-zero projection onto each motional mode direction, a single beam is sufficient to resolve multiple modes of different frequencies.

fig. 5.2, amplitude modulation of the RF trapping potential coherently excites secular motion if the DC and RF nulls do not coincide (that is, in the presence of a stray field $E \neq 0$). The equilibrium behaviour of the ion is then to oscillate at ω_{am}, with an amplitude proportional to the component of the stray field in the respective mode direction, and a well-defined phase relationship to the amplitude-modulation signal independent of the initial conditions, as we will show in the rest of this section.

To quantify amplitude and phase of the resulting motion, we record and digitally demodulate the arrival times of photons scattered during laser cooling, which is modulated through the Doppler effect, as illustrated in fig. 5.3. A single laser beam is sufficient for this as long as its wavevector $\boldsymbol{k}$ has a non-vanishing component along the direction of each mode.

The shift of the ion equilibrium position from the RF null in all three dimensions can thus be determined by successively modulating the potential at frequencies close to that of every mode of motion. The amplitude of motion correlated to the excitation signal depends linearly on the stray field along the direction of the addressed mode and crosses zero at the point where the DC and RF nulls coincide, providing a robust signal for a micromotion minimisation procedure.

5.2.1 *Equations of motion in the pseudopotential approximation*

Following the discussion in chapter 2, the time-dependent position x_i of the ion along one axis can be described by a Mathieu equation. Here, we will consider the case of a laser-cooled particle with an additional, unwanted electric field of uniform strength $\boldsymbol{E} = (E_i)_i$, which we wish to probe by amplitude-modulating the driving RF field at frequency ω_{am} and modulation depth h. We thus extend eq. (2.21) to the form (dropping tildes and mode

index subscripts)

$$\ddot{x} + 2\zeta\,\dot{x} + \left(\frac{\Omega_{\mathrm{rf}}}{2}\right)^2 \Big(a \, + 2q(1 + h\,\cos(\omega_{\mathrm{am}}t))\cos(\Omega_{\mathrm{rf}}t)\Big)\,x = \frac{E\,Ze}{m}, \tag{5.1}$$

where ζ gives the strength of linear viscous damping induced by laser cooling (see §5.2.4). Note that we have carried over the simplification to an uncoupled scalar equation from §2.2.1. As we are interested only in small modulation depths h, this is a reasonable approximation, but as discussed there, extensions of Floquet theory to the full matrix case could be employed for quantitative calculations [Hou08; Lan+12].

To understand the effect of the amplitude-modulated trapping potential on the ion's secular motion, it is instructive to consider eq. (5.1) in the pseudopotential approximation, where fast dynamics on the scale of Ω_{rf} are ignored. Here, for $|a| \ll |q| \ll 1$ and $\omega_{\mathrm{am}}, \zeta \ll \Omega_{\mathrm{rf}}$, we obtain[5]

$$\ddot{\bar{x}} + 2\zeta\dot{\bar{x}} + \left(\frac{\Omega_{\mathrm{rf}}}{2}\right)^2 \left(a + \frac{q^2}{2}(1 + h\cos(\omega_{\mathrm{am}}t))^2 \right)\bar{x} = \frac{E\,Ze}{m} \tag{5.2}$$

for the slowly-varying position $\bar{x}$. Defining

$$\omega_0 := \frac{\Omega_{\mathrm{rf}}}{2}\sqrt{a + \frac{q^2}{2}\left(1 + \frac{h^2}{2}\right)} \tag{5.3}$$

for the resulting frequency of secular motion and

$$\tilde{h} := \frac{h}{1 + h^2/2}\left(1 - \frac{a\Omega_{\mathrm{rf}}^2}{4\omega_0^2}\right) \tag{5.4}$$

as an "effective" modulation index, reflecting the fact that the contribution of Φ_{dc} to the radial curvature is static, we obtain

$$\ddot{\bar{x}} + 2\zeta\dot{\bar{x}} + \omega_0^2\left(1 + 2\tilde{h}\cos(\omega_{\mathrm{am}}t) + h\tilde{h}\cos(2\omega_{\mathrm{am}}t)\right)\bar{x} = \frac{E\,Ze}{m}. \tag{5.5}$$

Bringing this closer to the canonical form of a Mathieu equation, we can write

$$\ddot{\bar{x}} + 2\zeta\dot{\bar{x}} + \left(\frac{\omega_{\mathrm{am}}}{2}\right)^2 \left(A + 2Q\cos(\omega_{\mathrm{am}}t) + R\cos(2\omega_{\mathrm{am}}t)\right)\bar{x} = \frac{E\,Ze}{m}, \tag{5.6}$$

another damped Mathieu equation with stability parameters $A := (2\omega_0/\omega_{\mathrm{am}})^2$ and $Q := \tilde{h}A$ (up to the $\mathrm{O}(h^2)$ term oscillating at twice the modulation frequency, with amplitude $R := h\tilde{h}A$). Note that the parameter regime is quite different from the case describing ion motion in the radial plane, i.e. eq. (2.21), where $|a| \ll |q|$ is usually a small amount of anticonfinement resulting from the axially confining DC potential, and $|q| \ll 1$. Here,

---

[5]For instance by expressing x as a sum of two terms, $x = \bar{x} + \xi$, which are assumed to vary on different time scales.

we are instead interested in the regime close to the onset of parametric instability, that is, $A \approx n^2 > 0$ for the resonance index $n \in \mathbb{N}$, or $\omega_{\mathrm{am}} \approx 2\,\omega_0/n$.

Incidentally, note that eq. (5.6) has the same form as the equations of motion for a static harmonic potential superposed with an oscillating electric field linear in position, $E_x(t) = \omega_{\mathrm{am}}^2\, m\, Q\, x\, \cos(\omega_{\mathrm{am}} t)\,/\,(2\,Ze)$ (plus the small $\mathrm{O}(h^2)$ term). This could be a useful observation in applications where it might not be technically feasible to amplitude-modulate the trap RF drive signal at all, for instance because of an extremely narrow-band resonant RF matching circuit. In that case, the parametric excitation technique could still be used by instead applying the ω_{am} signal directly to the RF electrodes (e.g. combining it with the Ω_{rf} drive through a diplexer)[6]. In general, amplitude-modulation of the RF drive signal appears to be preferable, however, as the accuracy of the obtained minimum is not sensitive to a possible imbalance in electrode impedances between ω_{am} and Ω_{rf} (only between $\Omega_{\mathrm{rf}} \pm \omega_{\mathrm{am}}$), and the electrode voltages can more easily be filtered at frequencies near the secular modes of motion to avoid heating due to technical noise (which tends to be more prevalent at lower frequencies).

5.2.2 *Derivation of approximate solutions*

Solutions to eq. (5.6), that is, the ion trajectories resulting from parametric excitation, differ in character depending on the relative strength of excitation and damping. For the homogeneous case with no stray field, $E = 0$, the stability of the solutions can be investigated using standard techniques for the treatment of Mathieu equations. Analytical approximations for the stability boundaries can be obtained for small damping factors and detunings from resonance using a multiple-parameter perturbation method [WGW95]. For small amounts of damping, $\zeta \ll \omega_0$, the second parametric resonance is excited most strongly if $\omega_{\mathrm{am}} = \omega_0$, where the stability condition can be approximated to first order as

$$\tilde{h} \lesssim \frac{2\zeta}{\omega_0}. \tag{5.7}$$

In the stable region of parameter space, solutions for arbitrary initial conditions decay to the equilibrium at $\bar{x} = 0$ in the stable region of parameter space, at a characteristic rate given by the damping strength ζ. It can be shown [ST11] that the presence of an inhomogeneous term $E \neq 0$ does not alter the stability considerations. However, rather than tending towards a stationary equilibrium, solutions will then generally exhibit bounded steady-state behaviour. We have verified this to be an excellent approximation for the stability boundary for zero detuning in the experimentally relevant regime by numerical integration of eq. (5.1), although in practice nonlinearities in laser cooling and trap potentials significantly affect the transitory regime towards instability.

---

[6]Conversely, techniques where a modulating potential is directly applied to the trap RF electrodes, such as the quantum squeezing protocol demonstrated in ref. [Bur+19], should be accessible through amplitude-modulation sidebands near Ω_{rf} as well.

To now derive an approximation for bounded solutions for $A \approx 4$ (that is, $\omega_{am} \approx \omega_0$), we transition to normalised time $T := \frac{\omega_{am}}{2}t$, giving

$$\frac{d^2x}{dT^2} + 2B\frac{dx}{dT} + (A + 2Q\cos 2T + R\cos 4T)\,x = F, \tag{5.8}$$

for $B := \frac{2\zeta}{\omega_{am}}$, $F := \left(\frac{2}{\omega_{am}}\right)^2 \frac{EZe}{m}$ (and A, Q, R as in eq. (5.6)). All time-dependent factors in this linear differential equation are periodic in time with period π, so informed by a Floquet analysis of the system, we search for solutions of the form[7]

$$x(T) = \sum_{n=-\infty}^{\infty} c_n e^{2inT}, \tag{5.9}$$

with coefficients $(c_n)_{n\in\mathbb{Z}} \subset \mathbb{C}$ satisfying $c_{-n} = \overline{c_n}$ as $x(T) \in \mathbb{R}$. Inserting eq. (5.9) into eq. (5.8), i.e. following a harmonic balance approach, yields a system of recurrence relation equations for the coefficients:

$$2Q(c_{n-1} + c_{n+1}) + R(c_{n-2} + c_{n+2}) + 2(A - 4n^2 + 4niB)c_n = \delta_{n,0}2F. \tag{5.10}$$

The magnitude of c_n will fall rapidly with $|n|$ due to the $4n^2$ factor, as Q, R and B are small for typical conditions. We can thus obtain an analytical solution by truncating the system at an appropriate $|n|$ for the desired degree of accuracy.

To do so for small excitation amplitudes, we substitute $Q = \tilde{h}A$ and $R = \alpha\tilde{h}^2A$ (with $\alpha := h/\tilde{h} \approx 1$), and expand to second order in $\tilde{h}$ (for which including terms up to $|n| = 2$ is sufficient), obtaining

$$c_0 = \frac{F}{A}\left(1 - \tilde{h}^2\frac{2A(4-A)}{(4-A)^2 + (4B)^2}\right) + O(\tilde{h}^3),$$

$$c_1 = \tilde{h}\frac{F}{4 - A - 4iB} + O(\tilde{h}^3), \tag{5.11}$$

$$c_2 = \tilde{h}^2\frac{F(2A + \alpha(4 - A - 4iB))}{2(4 - A - 4iB)(16 - A - 8iB)} + O(\tilde{h}^3).$$

It is worth noting that the ansatz (5.9) only covers bounded solutions, and only the case of excitation near integer subharmonics of the motional frequency (of which $\omega_{am} \approx \omega_0$ is the strongest). Parametric resonance near $2\omega_0$ and other subharmonics thereof excites motion at half the modulation frequency, which is consequently not captured by the above solution.

---

[7]With thanks to P. Drmota for the suggestion of introducing complex coefficients here instead of separate cosine and sine terms, leading to a pleasingly compact recurrence relation.

5.2.3 Behaviour of the first-order steady-state solution

Considering the expressions from eq. (5.11) to first order in $\tilde{h}$ ($\tilde{h}^2 \ll 1$), all coefficients c_n with $|n| > 1$ vanish, and substituting back physical units, we obtain the average ion position in

$$c_0 = \frac{EZe}{m\,\omega_0^2} \tag{5.12}$$

as expected for the equilibrium position given a constant force from the external electric field E. Introducing the normalised excitation frequency and damping parameters

$$f := \frac{\omega_{\mathrm{am}}}{\omega_0}, \quad d := \frac{\zeta}{\omega_0}, \tag{5.13}$$

and defining $C := 2c_{-1}$ for compactness of notation such that $\bar{x}(t) = c_0 + \mathrm{re}(C\,e^{-i\omega_{\mathrm{am}}t}) + O(\tilde{h}^2)$, we obtain

$$C = \frac{2\tilde{h}}{f^2 + 2ifd - 1}c_0, \tag{5.14}$$

which is which is valid to second order in $\tilde{h}$, or in terms of magnitude and phase,

$$|C| = \frac{2\tilde{h}}{\sqrt{(f^2 - 1)^2 + 4f^2d^2}}|c_0|, \quad \begin{aligned}\arg(C) &= \arg(E(f^2 - 2ifd - 1)) \\ &= \mathrm{Atan}(-2fdE, (f^2 - 1)E)\end{aligned}. \tag{5.15}$$

In the limit of small modulation depth, the system thus behaves like a damped, driven harmonic oscillator. If the mode is excited exactly on resonance, the resulting motion is phase-shifted by $\pm\pi/2$ compared to the applied amplitude modulation (depending on the sign of E). At modulation frequencies much smaller or much larger than the mode frequency the shift is 0 or π, respectively, with the width of the transition region being proportional to the damping coefficient.

Figure 5.4 shows the behaviour predicted by eq. (5.15) for a set of parameters typical for a microfabricated surface trap, along with numerical simulations for the full motional dynamics as per eq. (5.1). For the small modulation depth of $h = 0.007$ used here, the approximation from eq. (5.15), which includes only terms linear in h, only slightly underestimates the amplitude near resonance.

Crucially, the correlated amplitude C is linear in c_0, and thus in the stray field strength E. This linear dependence enables a simple stray field compensation procedure: for an arbitrary fixed modulation frequency ω_{am} close to the frequency of one mode of motion, measure C for a variety of static compensation fields; its zero-crossing then marks the compensated point where the stray field is exactly cancelled.

For any given set of system parameters, excitation exactly at resonance ($f = 1$) yields the maximum signal $|C|$, so in principle this would be the optimal working point for the calibration procedure. However, at this point, the observed amplitude C is maximally sensitive to small deviations in the mode frequency, whether caused by trap anharmonicity
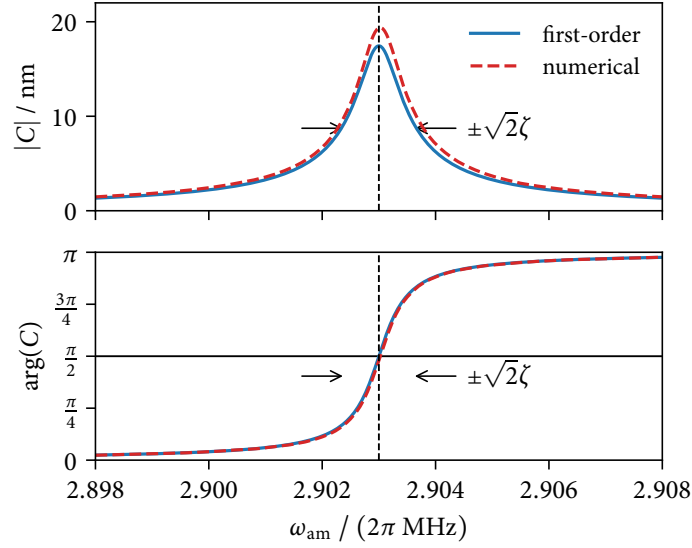
Figure 5.4: Predicted amplitude and phase of steady-state oscillation following excitation at ω_{am} with modulation index $h = 0.007$ and damping $\zeta = 2.4\,\text{ms}^{-1}$ in the presence of a stray field with $EZe/m = -1\,\text{m/s}^2$, for a motional mode described by $a = -0.00259$, $q = 0.179$, and $\Omega_{\text{rf}} = 2\pi \times 50\,\text{MHz}$ ($\omega_0 = 2\pi \times 2.903\,\text{MHz}$). The solid blue curves show the result obtained to first order in h in the pseudopotential approximation, given by eq. (5.15). The dashed red curves were obtained by numerically integrating the full equations of motion. For ease of comparison, the latter were shifted horizontally by the $2\pi \times 17.51\,\text{kHz}$ difference between the mode frequency obtained from full simulations without modulation and that in the pseudopotential approximation, $\omega_0 \approx \Omega_{\text{rf}}/2\,\sqrt{a + q^2/2}$ (ω_0 is indicated by the vertical dashed lines).

(changes in local pseudopotential curvature as the ion position is shifted by different static compensation fields), or for instance by technical drifts in the RF voltage level across the duration of the calibration procedure. Concretely, consider the relative change in signal caused by a small shift in the actual mode frequency,

$$\left| \omega_0 \frac{\partial C}{\partial \omega_0} \middle/ C \right| = \frac{2}{\sqrt{(f^2 - 1)^2 + 4f^2 d^2}} \quad \left(= \frac{|C|}{\tilde{h}|c_0|} \right), \tag{5.16}$$

which is in fact maximal at resonance and decreases with increased detuning exactly as quickly as the magnitude of the signal $|C|$ does. Thus, detuning ω_{am} by a small amount chosen based on the observed stability of ω_0 can greatly increase the robustness of the procedure in practice, at the cost of a reduction in signal strength. To some extent, the loss in signal can be counteracted by an increase in modulation depth, though eventually the approximations made here ($h \ll 1$) will break down.

5.2.4 Approximate ion dynamics in the linearised two-level picture

To connect the results obtained in the last section to the experiment, it is necessary to model the internal degrees of freedom of the trapped ion, both in order to extract the damping coefficient ζ resulting from laser cooling, and to derive the modulation in fluorescence count rate used to infer the motional amplitude. A simple model adequate for our

purposes is that of a two-level system always assumed to be in equilibrium (that is, with internal state dynamics fast compared to ω_0).

Linear damping through Doppler cooling. Consider a two-level system with ground state g and excited state e, connected by a transition with linewidth Γ, and driven to equilibrium by an external field (here, the cooling laser) with Rabi frequency Ω_{eg} and a detuning Δ from resonance, which scatters photons at a rate R given by eq. (2.60). Due to the Doppler shift, ion motion gives rise to an instantaneous detuning $-k\dot{x}$, where k is the projection of the wavevector of the laser beam onto the mode direction. Making the simplifying assumption that the ion velocity changes slowly [Lei+03], such that the internal state is in equilibrium at any point, we can write

$$R = \Gamma \left(\rho_{ee}(\Delta) - \rho'_{ee}(\Delta)\, k\dot{x} + O\big((k\dot{x}/\Gamma)^2\big) \right). \tag{5.17}$$

Neglecting momentum diffusion from the isotropically emitted photons as well as any temporal correlations between absorption and emission events, laser cooling can then be modelled through the average scattering force $F = \hbar k R$, giving the linear damping term as

$$\zeta = -\frac{1}{2m} \frac{\partial F}{\partial \dot{x}}. \tag{5.18}$$

Motion-dependent fluorescence modulation. The Doppler shift also leads to a modulation of the scattering rate dependent on the ion motion. If a fraction of the emitted photons are registered by some detection system with overall efficiency η, and the ion velocity is given by $\dot{x}(t) = \omega_{\mathrm{am}} \operatorname{im}(C e^{-i\omega_{\mathrm{am}}t})$ for the complex amplitude C defined in §5.2.3, the observed count rate is given by

$$r(t) = \eta R(t) \approx \eta \left(R_0 + \frac{4k\omega_{\mathrm{am}}}{\Gamma} \frac{s\Delta}{\left(1 + s + \left(\frac{2\Delta}{\Gamma}\right)^2\right)^2} \operatorname{im}(C e^{-i\omega_{\mathrm{am}}t}) \right), \tag{5.19}$$

where $R_0 = R(\dot{x} = 0)$ is the average scattering rate which coincides with the unexcited case.

As such, by measuring the frequency component of the count rate at the excitation frequency, which we call the correlated rate S, we can recover the excitation amplitude C via

$$S := \left\langle r(t) e^{i\omega_{\mathrm{am}}t} \right\rangle_t = -i \frac{2\eta k\omega_{\mathrm{am}}}{\Gamma} \frac{s\Delta}{\left(1 + s + \left(\frac{2\Delta}{\Gamma}\right)^2\right)^2} C, \tag{5.20}$$

where $\langle \cdot \rangle_t$ denotes the average over an integer multiple of excitation periods.

In the experiment, S is easily estimated by recording the arrival times $(t_n)_{n=1,\dots,N}$ of N photons over some time $T \gg 2\pi/\omega_{\mathrm{am}}$ and digitally demodulating the signal as

$$\langle r(t)\mathrm{e}^{\mathrm{i}\omega_{\mathrm{am}}t}\rangle_t \approx \frac{1}{T}\sum_{n=1}^{N}\mathrm{e}^{\mathrm{i}\omega_{\mathrm{am}}t_n}. \tag{5.21}$$

The finite number N of observed photons leads to a statistical variance, i.e. photon shot noise, of $N/(2T^2) = \eta R_0/(2T)$ per complex component of S. S is linear in the motional amplitude C and thus, as shown in the previous section, in the stray electric field E. If we use this procedure to estimate E, the statistical variance for an observation time T is

$$\sigma_E{}^2 = \underbrace{\left|\frac{E}{S}\right|^2 \frac{\eta R_0}{2}}_{=:\alpha^2} \frac{1}{T}, \tag{5.22}$$

where we have introduced the overall (inverse) sensitivity α of the stray field measurement (unit: $\mathrm{Vm}^{-1}/\sqrt{\mathrm{Hz}}$). Substituting the concrete expressions for S and C obtained earlier, we thus obtain the sensitivity of the method in this linearised model as

$$\alpha^2 := \left(\frac{m}{2\tilde{h}Ze}\right)^2 ((f^2-1)^2 + 4f^2d^2)\frac{\Gamma^3\left(1+s+\left(\frac{2\Delta}{\Gamma}\right)^2\right)^3}{16\eta k^2\Delta^2 s}, \tag{5.23}$$

noting that the laser parameters s and Δ also enter the normalised damping parameter d (see eqs. (2.60), (5.13) and (5.18)).

Choosing parameters for sensitivity. Without further constraints, attempts to choose experimental parameters for optimal sensitivity by minimising eq. (5.23) are ill-fated. For excitation exactly at the mode frequency ($f = 1$), reducing the damping coefficient would lead to arbitrarily large ion orbits and thus relative signal rates $S/(\eta R_0)$. In reality, however, nonlinearities of trap and cooling forces [Ake+10] and the onset of parametric instability limit the achievable amplitudes. The simple, linear model derived here is thus inadequate to derive an ultimate bound on the sensitivity of this method. Nevertheless, if the allowable motional amplitude $|C|$ is fixed (for instance, as limited by trap uniformity), eq. (5.23) attains a minimum at $-\Delta = \Omega_{eg} = \Gamma/2$. Similarly, if $\tilde{h}$ and f are fixed (choosing, for instance, a detuning $1 - f \approx 5 \cdot 10^{-3}$ to decrease sensitivity to trap frequency drifts), the sensitivity is again found to be maximal near $\Delta \approx -\Gamma/2$ over a range of mode frequencies and linewidths typical of trapped ion experiments. This suggests $\Delta = -\Gamma/2$ to be a suitable starting point for further empirical optimisation.

5.3 *Experimental parametric excitation response*

We demonstrate this method using a single $^{88}\mathrm{Sr}^+$ ion in the "Alice" node described in chapter 3. For this, we operate the trap with electrode voltages chosen to result in typical mode
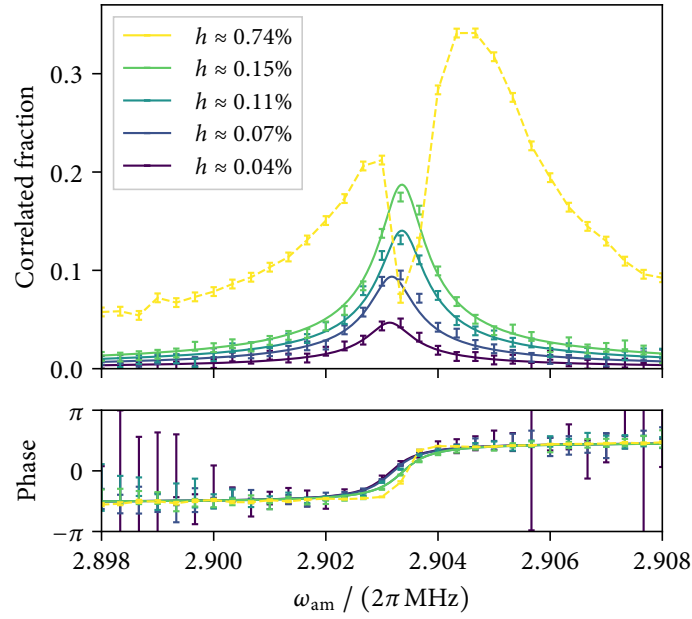
Figure 5.5: Observed amplitude and phase of the correlated photon count rate versus the modulation frequency. The signal amplitude S is given as a fraction of the approximately constant total count rate of $\eta R_0 = 15\,\mathrm{kHz}$, for cooling parameters $\Delta = -2\pi \times 13\,\mathrm{MHz}$, $s \approx 0.9$. The excitation frequencies are near the higher-frequency radial mode, and we intentionally apply an offset field of 5 V/m along the mode direction. Data for five different modulation indices h are shown: For the smaller four, the solid lines lines show a fit of the model from eq. (5.15), allowing for an individual offset in mode frequency and a constant phase shift common to all curves (yielding a damping rate $\zeta = 2.4\,\mathrm{ms}^{-1}$). For excitation at $h \approx 0.74\,\%$, where nonlinear effects dominate, the dashed line connects the points to guide the eye. Error bars show 1σ statistical errors from photon shot noise (50 000 photons per point).

frequencies of 1.85 MHz, 2.28 MHz, and 2.90 MHz for the axial and two radial modes, for which a static quadrupole field term is added such that the radial mode vectors are rotated by 28° about the trap axis (resulting in overall angles of 70° and 51° between the cooling beams and the low- and high-frequency radial modes, respectively). The amplitude modulation signal is generated using a DDS which is phase-referenced to the same clock used for photon timestamping, as described in §3.2.1, which enables a purely digital demodulation of the fluorescence rate signal.

Figure 5.5 shows amplitude and phase of the correlated scattering rate signal S as a fraction of the total fluorescence rate for excitation near the higher-frequency radial mode, with the ion shifted from the RF null by 5 V m^{-1} along the direction of that mode. For these data, the ion is continuously cooled. After switching on the amplitude-modulation signal, a 300 µs settling period reduces the influence of the initial conditions. Timestamps are then recorded for each PMT click until a pre-defined number of photons (here 50 000) has been observed. The signal is demodulated according to eq. (5.21) by comparing the
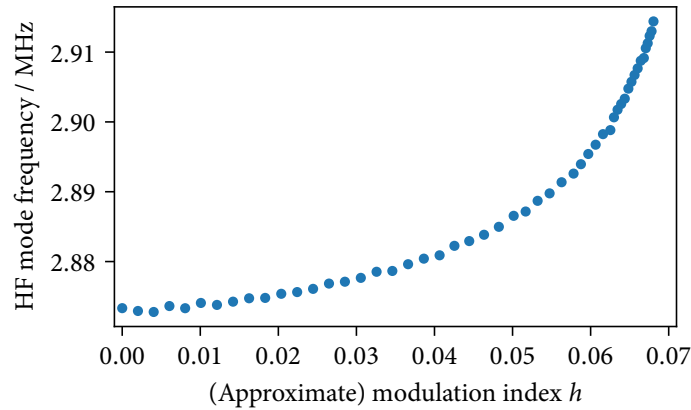
Figure 5.6: Shift of the HF mode frequency versus amplitude of an off-resonant amplitude modulation tone at $\omega_{am} = 2\pi \times 2.5\,\text{MHz}$, as observed using a "tickle" measurement. The statistical uncertainty of each point is smaller than the marker size, but the accuracy is affected by kHz-level drifts in the unperturbed mode frequency over the duration of the measurement (points acquired in random order).

timestamps to the phase origin of the parametric excitation envelope signal[8].

The behaviour for a number of different excitation amplitudes is shown, with cooling parameters $\Delta = -2\pi \times 13\,\text{MHz}$, $s \approx 0.9$. The absolute scaling of the depth of the applied modulation after the trap resonator was calibrated to about one significant figure by probing the signal at a capacitive voltage divider after the resonant circuit with an oscilloscope. For small modulation indices, the behaviour agrees with the linear model derived in §5.2.3, shown in the form of fit curves.

The observed rate grows linearly with modulation amplitude (and distance from the RF null, which is constant here) until non-linearities in trap and motional coupling start to become significant. Empirically, we observe a gradual saturation behaviour towards a correlated fraction of 0.35 of the total fluorescence rate, i.e. a scattering rate amplitude modulation depth of 0.7. For larger excitation amplitudes, a sharp drop in correlated photon rate is observed near resonance. This effect cannot be reproduced by the linear response model; we attribute it to ion response nonlinearities (when $|k\dot{x}|$ reaches a significant fraction of the linewidth), trap nonlinearities (for large orbits), and the onset of parametric instability. Empirically, we observe that for excitation near ω_0 the ion remains confined in the trap even under those conditions. Higher excitation amplitudes at larger detunings can be a useful working point for the stray field compensation algorithm discussed in the following section, as the effectively broadened response reduces the sensitivity to changes in mode frequency.

Another effect visible in fig. 5.5 not captured by the model is a systematic increase in observed mode frequencies towards higher excitation amplitudes (ω_0 was allowed to float for each curve in the joint fit). The origin of this effect is unclear. In the pseudopotential

---

[8]There is a non-zero total delay through the RF modulation and detection chain, so a linear phase equivalent to 1.8 µs was subtracted from the data for display purposes, recovering an asymptotically flat phase response far off resonance. This does not affect the stray field compensation algorithm described in §5.4.

approximation, a (larger) increase in mode frequency with increasing modulation index h is expected according to eq. (5.3). On the other hand, numerical integration of eq. (5.1), the full equation of motion, for parameters $a = -0.0026$ and $q = 0.17$ suggests a decrease in observed mode frequency for pure amplitude modulation instead. To probe this effect in isolation, trap RF modulation was applied at $\omega_{am} = 2\pi \times 2.5$ MHz, far from resonance with either motional mode. The HF mode frequency was then determined using the "tickle" method, where a resonant electric field applied to one of the trap electrodes excites the ion motion. We detect the change in total fluorescence rate due to the large motional amplitudes produced this way. An increase of mode frequency with modulation index is indeed observed, as shown in fig. 5.6, but matches neither the result from the pseudopotential approximation nor the numerical simulations of the full one-dimensional model from eq. (5.1)[9]. In practice, this shift is not a hindrance for use of the stray field compensation method as the mode frequency can simply be determined at the target modulation index. However, it needs to be accounted for when deriving the amplitude-modulation settings from mode frequencies calibrated using a different method and large modulation depths are to be used.

To determine the optimum compensation fields, the modulation frequency ω_{am} is kept fixed, but a varying offset field E is applied instead, as shown in fig. 5.7. Matching the prediction from eq. (5.15) and eq. (5.20), the correlated rate is observed to change linearly with the applied electric field. The setting that minimises the micromotion amplitude is thus easy to determine by linear regression, which is computationally cheap and robust to statistical perturbations.

5.4   Multi-dimensional stray field compensation

By probing for parametric excitation near each of the mode frequencies, we can measure the stray field component along each mode. Thus, stray fields in all spatial directions could be compensated. As in most linear Paul traps, the RF potential curvature along the trap axis is very small in our system by design, so we present data only for the plane spanned by the two radial modes in the following.

---

[9]One possible origin for this discrepancy lies in the static DC quadrupole field applied to rotate the mode directions away from the trap plane, which leads to the treatment of the equations of motion as three uncoupled scalar equations not to be strictly accurate (see the discussion in §2.2.1/eq. (2.21)). Whether this plays a role or e.g. effects from non-uniform curvature dominate would be straightforward to investigate by simulating the full three-dimensional dynamics of the ion in the fields derived from the model of the trap geometry. Whatever the explanation, any such effects should not appreciably offset the excitation minimum extracted for low modulation strengths, however. In this context, it is still somewhat unfortunate that the amplitude-dependent mode frequency shift is not modelled well by the analytical or one-dimensional numerical simulations, as it could have otherwise served as a convenient calibration of the actual modulation index after the band-pass filter formed by the resonant trap RF matching circuit.
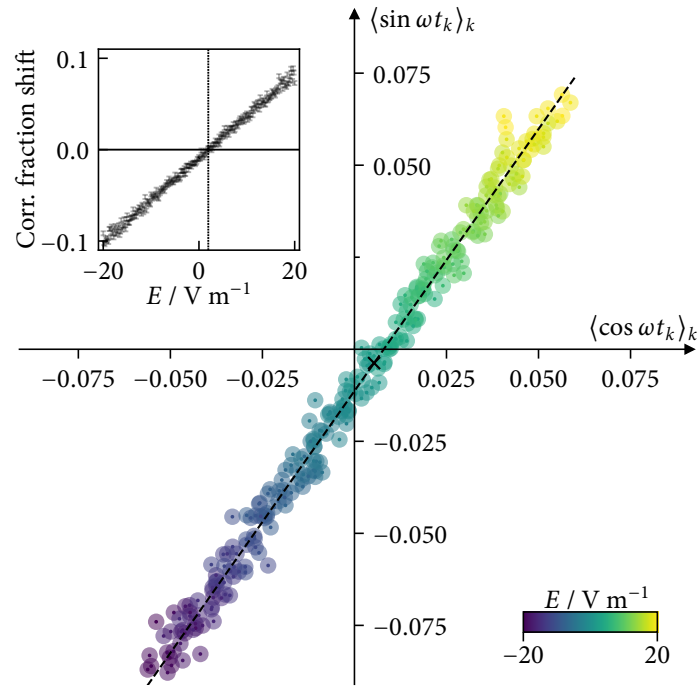
Figure 5.7: Correlated scattering rate signal for excitation close to a 2.88 MHz radial mode (detuning $-22\,\mathrm{kHz}$, $h \approx 0.02$), given as fraction of the 30 kHz total count rate for varying compensation fields along the mode direction. The dots are coloured to denote the applied fields, ranging from $-20\,\mathrm{V\,m^{-1}}$ (dark) to $20\,\mathrm{V\,m^{-1}}$ (light) relative to an arbitrary value close to the compensation optimum; their size corresponds to the 1σ statistical error from 100 000 observed photons each. The dashed line shows a linear complex least-squares fit to the data as a function of E-field. In the inset the same data is shown in one-dimensional form, with the vertical axis now giving the projection of each point onto the best-fit line, where zero corresponds to the point on the line closest to the coordinate origin. The symbols show the statistical error, confirming that the magnitude of excitation is linear to excellent approximation. The cross in the main plot denotes the inferred compensation optimum, that is, the point on the fit line closest to the origin. The same point, corresponding to $E = 1.95(3)\,\mathrm{V\,m^{-1}}$, is indicated in the inset by the dotted line.

As illustrated in fig. 5.8, this is a multi-dimensional optimisation problem. Even when only varying the applied compensation fields along a single direction, there will usually be a point for each mode where the excitation vanishes, i.e. the shift between DC and RF minimum along that mode direction is zero. However, unless the spatial orientation of the motional modes is known to high precision so that an appropriate compensation voltage basis can be chosen (taking into account, for instance, the contributions of any stray quadrupole potentials), such a one-dimensional search will in general not yield enough information to determine the voltage set that positions the ion exactly on the RF null.

This is further illustrated in fig. 5.9, which shows the motional amplitudes observed as the compensation fields vary in both radial directions for a fixed excitation frequency. For each mode, there exists a linear minimum in excitation amplitude, where the projection of stray field onto the mode direction vanishes, and the ion is merely translated perpendicular to the mode direction. A linear fit to each dataset reveals mode directions of $-63°$ and $26°$ relative to the trap plane (based on the geometry of the compensation electrodes), in
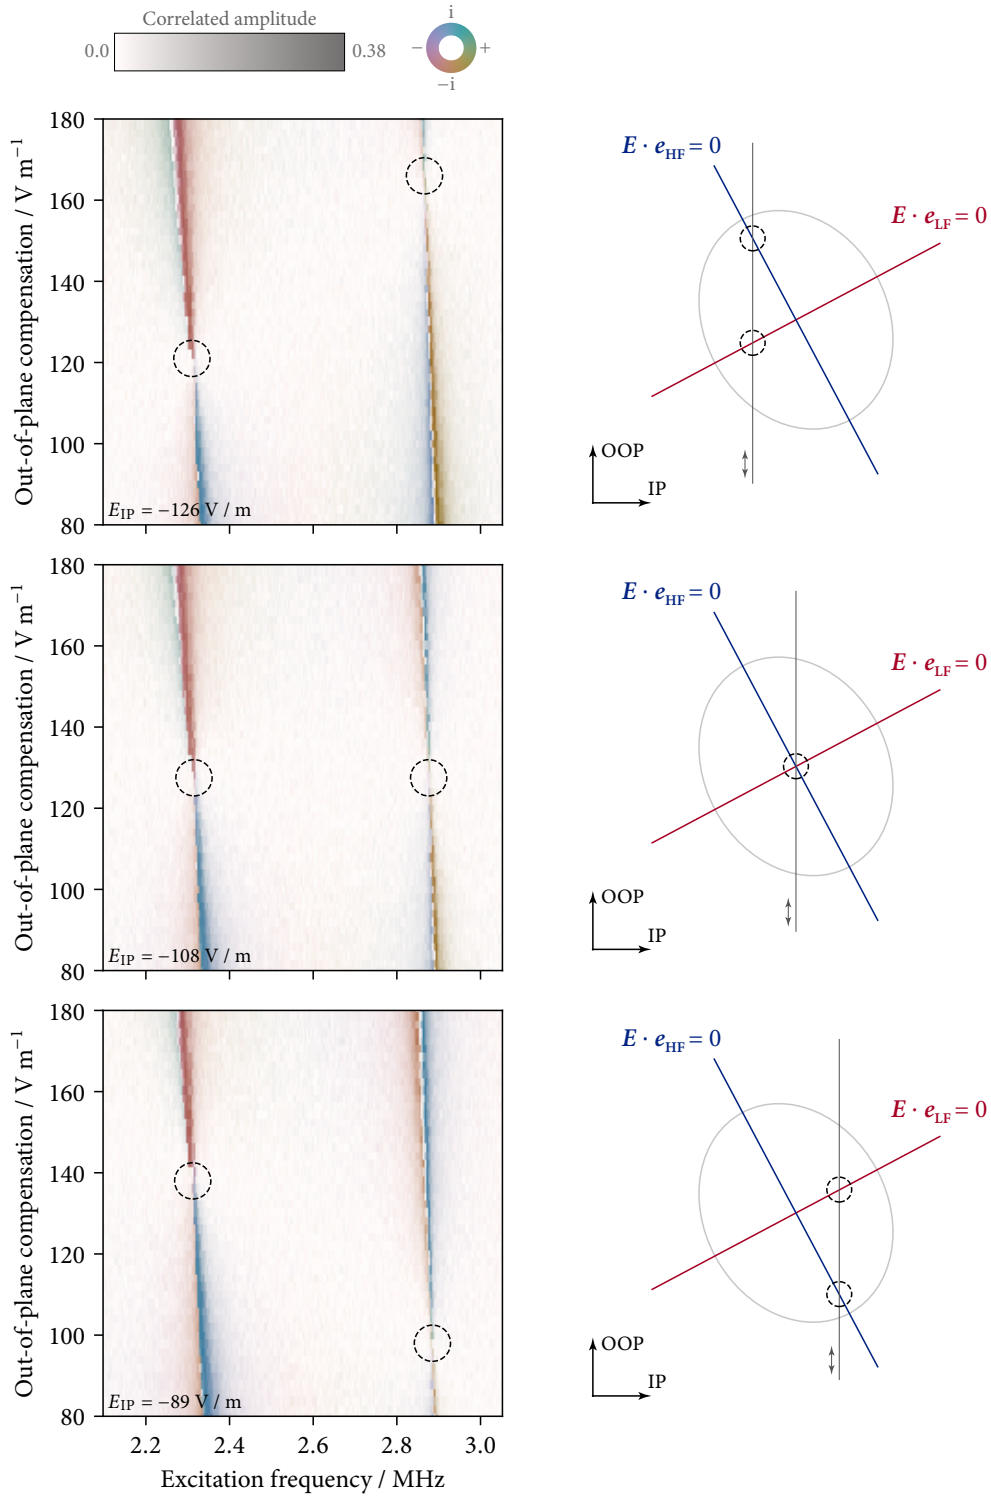
Figure 5.8: Complex fractional amplitude of correlated photon arrival times, as a function of excitation frequency and applied field normal to the trap plane E_{OOP}, for three different in-plane compensation fields E_{IP} (as illustrated in the radial plane sections to the right, where e_{LF} and e_{HF} refer to the directions of the lower- and higher-frequency radial modes). Colour shows phase, and brightness shows amplitude (extracted from 10 000 photons observed per point). The dashed circles highlight the points where the projection of micromotion onto the respective radial mode vanishes. Only when E_{IP} is chosen such that the stray field in the in-plane direction is minimised (centre plot) can both modes be compensated. The dependence of radial mode frequencies on ion position due to trap anharmonicity is visible as a slant to the resonance lines.
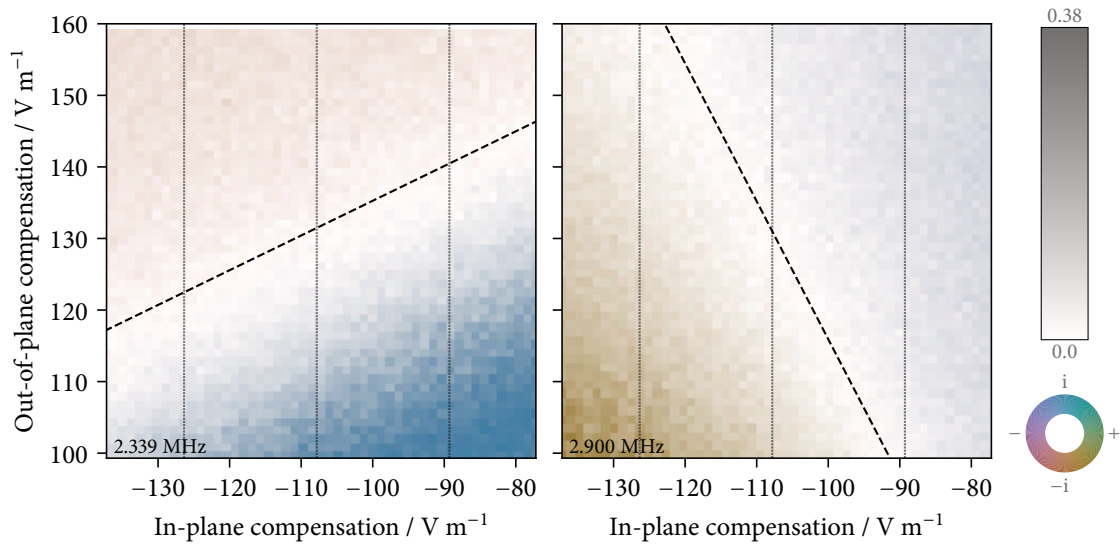
Figure 5.9: Stray field compensation in two dimensions. The plots show the complex fractional amplitude of photon arrival times correlated with the fixed excitation frequency of 2.339 MHz and 2.900 MHz, respectively, for a detuning of 22 kHz relative to each radial mode in the scan centre. Along the dashed lines, the residual displacement of in the direction of the given mode is minimised; all radial micromotion is compensated in the centre of the plot, where both lines would intersect. The dotted vertical lines mark the values for E_{IP} used in the three scans from fig. 5.8 for visual comparison. Both mode frequencies increase towards negative out-of-plane fields, reducing the detuning and giving rise to the stronger signals towards the lower edge of the plots.
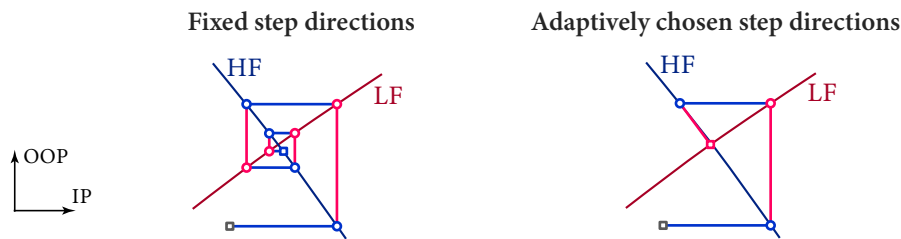
Figure 5.10: To compensate simultaneously the stray fields in both directions, a fixed-point iteration scheme can be employed, where the excitation on either mode is minimised in an alternating fashion by applying compensation fields in some pre-defined basis (here in the IP and OOP directions). This simple strategy is successful as long as the mode orientation is known to better than 45° (left), but an adaptive algorithm, which at each point estimates the local mode directions from previous measurements to avoid moving the ion off the other null, can greatly speed up convergence (right). (On the approximately straight lines in the background, no motion is excited in the LF or HF modes, with the place where they meet being the compensation optimum. The markers, connected by lines, denote the result from each iteration step of the respective compensation algorithm, starting from an arbitrarily chosen point.)

reasonable agreement with the 28° intended mode axis rotation angle. At the point where both lines intersect, the trap is compensated: the ion is centred on the RF null and no parametric excitation is observed.

5.4.1   Iterative compensation algorithm

Acquisition of two-dimensional data of this nature is unnecessarily time-consuming. As illustrated in fig. 5.9 (c), as long as the direction of the applied compensation fields relative to the normal modes is known to better than 45°, a simple fixed-point iteration method, where excitation for each mode is minimised in an alternating fashion, will converge to the compensation optimum. The orientation of each mode, expressed in terms of compensation voltages, is typically constant to good approximation, at least close to the compensated point. If, at each step, the direction of the currently inactive mode(s) is estimated from previous iterations, and the probe direction is then chosen such as not to displace the ion in that direction, convergence can be greatly accelerated. Note that the obtained motional mode directions are not necessarily normal when expressed in the compensation voltage basis, for instance if the fields generated by each electrode voltage set are not in fact orthogonal at the ion position, or of different magnitude.

In practice, a further complication arises from the fact that the observed mode frequencies will shift slightly as a function of the applied compensation voltages. One reason for this is that the DC and RF potential curvatures are not necessarily constant over the volume probed by the ion positions, especially in surface electrode traps. The potentials introduced by the compensation electrodes are also not in general curvature-free. In fig. 5.8, this shift appears as a slant to the resonance features. To some extent, this can be mitigated by choosing a parametric excitation frequency far enough detuned from resonance that the change in response remains acceptably small over the region of interest. However, if the region of compensation fields to be investigated is large, the detuning required for this might be too big to be practical, given that the achievable damping rate and amplitude modulation index will be limited.

To robustly find a set of compensation electrode voltages without prior knowledge of the mode structure, we thus employ a multi-step algorithm (still alternating between modes). Given a frequency estimate for the modes of interest and an initial guess and search radius for the respective compensation voltage sets to apply, iterate over the following steps:

1. Determine the frequency of the mode in question at the current compensation settings plus a shift of one search radius in one direction. This can for instance be achieved by varying the excitation frequency and fitting a step of π in phase to the observed correlated amplitude (see fig. 5.5).

2. Repeat step 1 for voltages shifted by one radius in the opposite direction, thus obtaining a linear model for mode frequency as a function of compensation voltage.

3. Then, vary the applied compensation field over the search range while parametrically exciting at some fixed frequency offset to the approximation just determined, and
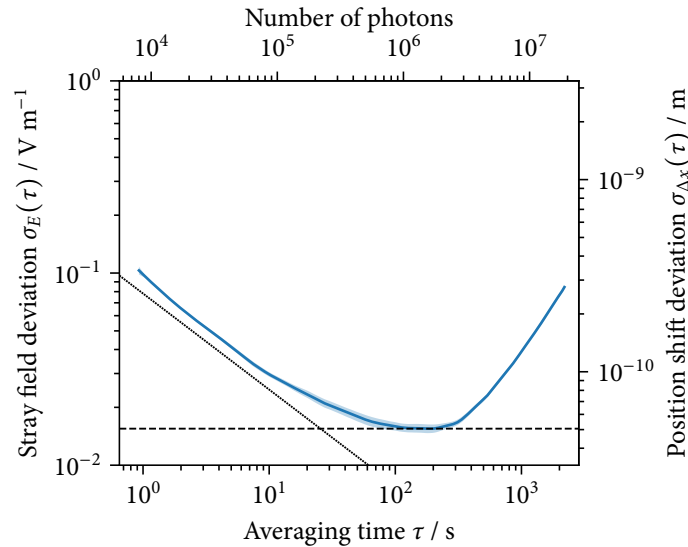
Figure 5.11: Stability evaluation of the compensation field in the direction of the high-frequency mode. Shown is the overlapping two-sample (Allan) deviation $\sigma_E(\tau)$ for the inferred compensation field for different averaging times τ, converted into a positional uncertainty for the 2.92 MHz ^{88}Sr$^+$ mode on the right axis. The shaded region indicates the statistical uncertainty (1σ) due to photon shot noise, obtained from parametric bootstrapping. The dotted line denotes the error expected purely from photon shot noise at the fitted sensitivity, and the dashed line indicates the minimum uncertainty of $\sigma_E = 0.015\,\text{V m}^{-1}$ achieved at an averaging time of $\tau \approx 150\,\text{s}$.

determine the point of minimal excitation by a complex linear fit to the count rate correlations (see fig. 5.7).

Steps 1–3 are applied, alternating between all motional modes, until the change in estimate between iterations is smaller than some pre-set target precision. After having cycled through each mode two times the mode direction estimates can be refined, allowing compensation voltage sets to be chosen adaptively such as to minimally disturb the result of previous iterations. As the RF null is approached, the excitation amplitude can be increased while decreasing the scan radius to improve the sensitivity. The mode frequency calibrations, i.e. steps 1–2, can be skipped if the mode parameters are well-known, for instance when repeatedly running calibrations to track small stray field changes while the setup remains nominally unchanged.

5.5 Statistical uncertainty

The main source of statistical measurement uncertainty is photon shot noise, i.e. the Poissonian distribution of observed photon counts. As such, large light collection efficiencies are desirable to reduce data acquisition time for a given target precision. Beyond that, there is a trade-off between sensitivity and robustness: For operation in automated calibration routines, we typically choose larger damping rates and detunings ($\sim 10\,\text{kHz}$) for robustness

against slow radial mode frequency drifts (which we tend not to track with very high resolution) as well as lower excitation amplitudes to handle larger stray field changes without leaving the linear region. With these conservative settings, stray fields can be compensated to well below 1 V/m from 200 000 photons per direction in our apparatus.

When choosing larger excitation amplitudes and smaller detunings ($\sim 3\,\text{kHz}$), sensitivities of the correlated fraction to stray fields of $|S|/E \approx 1.4\,\text{kHz}/(\text{V m}^{-1})$ at total count rates of $\eta R_0 = 17\,\text{kHz}$ can routinely be reached in our apparatus, giving a statistical noise floor of $\alpha = 0.1\,\text{Vm}^{-1}/\sqrt{\text{Hz}}$ per mode direction.

To estimate the practical resolution limit achievable with our method given the slow, technical drifts in the voltage sources and the environment (e.g. due to temperature changes), we consider Allan's two-sample variance for the inferred stray electric field. We continuously measure the correlated amplitude for a range of applied compensation fields chosen uniformly between $\pm 2.5\,\text{V m}^{-1}$, for excitation with $h \approx 0.03$ detuned $-3\,\text{kHz}$ from the high-frequency radial mode, cooling parameters $s \approx 1$, $\Delta = -13\,\text{MHz}$, and $N = 1000$ photons per compensation field setting.

Let $\bar{E}_{i;l}$ denote the estimate for the compensation field that minimizes the excess micromotion amplitude as determined by linear regression (as discussed in §5.4; see fig. 5.7) over the subset of correlated amplitude measurements from index i to $i + l - 1$, which we assume to originate from $L - 1$ equal time intervals $\Delta t = \frac{N}{\eta R_0}$. For technical simplicity, we acquire an equal number of photons per point, making this only true in approximation due to the Poissonian count statistics. The overlapping two-sample variance $\sigma_E(\tau)^2$ for an averaging period $\tau = l\,\Delta t, l \in \mathbb{N}$, is then

$$\sigma_E^2(l\,\Delta t) = \frac{1}{2(L - 2l + 1)} \sum_{i=1}^{L-2l} \left(\bar{E}_{i;l} - \bar{E}_{i+l;l}\right)^2. \tag{5.24}$$

Figure 5.11 shows the dependence of this variance on the averaging time for a representative data set, for averaging windows of $l = 8$ to $19\,640$. At short times, the variance drops with the square root of the averaging time, reflecting the floor set by the photon shot noise. Around $\tau \approx 1\,\text{min}$, a plateau is reached, where further averaging does not significantly improve the variance. The achieved minimal uncertainty is $\sigma_E = 0.015\,\text{V m}^{-1}$, which, given the 2.92 MHz frequency of the mode used for this dataset and the charge-to-mass ratio of $^{88}\text{Sr}^+$, corresponds to an uncertainty in position relative to the RF null of $\sigma_{\Delta x} = 34\,\text{pm}$, or an equivalent force of $\sigma_F = 1.7\,\text{zN}$. At longer averaging times, slow drifts in the stray fields or electrode voltages dominate.

5.6   *Systematic effects*

The data in the previous section only covers the statistical uncertainty in the compensation fields. To check for any unknown systematic effects, we compare the results against

two phase-modulation-based methods, albeit only along the one direction parallel to the trap plane in which the relevant laser beams propagate. Here we find small but statistically significant differences between the three approaches: At one point, a typical set of compensation fields obtained by the parametric excitation method was $E_{\mathrm{IP}} = -119.6(1)\,\mathrm{V\,m^{-1}}$ and $E_{\mathrm{OOP}} = 147.8(1)\,\mathrm{V\,m^{-1}}$. Subsequently adjusting E_{IP} to instead minimise the correlations between 422 nm photons scattered during Doppler cooling and the trap RF drive yielded a E_{IP} value lower by $3.3(1)\,\mathrm{V\,m^{-1}}$. As a third method, we can minimise the Rabi frequency on the first negative micromotion sideband of the narrow-linewidth 674 nm $\left|S_{1/2}, m = -1/2\right\rangle \leftrightarrow \left|D_{5/2}, m = -3/2\right\rangle$ transition: this minimum occurred at an E_{IP} value lower by $7.2(1)\,\mathrm{V\,m^{-1}}$ compared with that given by the parametric excitation method. Some weeks prior to that, the same 674 nm sideband method had given a result higher by $1.1(1)\,\mathrm{V\,m^{-1}}$.

Taking some known, but relatively obscure systematic shifts affecting the other methods into account, the differences for the 422 nm RF correlation and 674 nm sideband methods reduced to $-1.0(1)\,\mathrm{V\,m^{-1}}$ and $-3.7(1)\,\mathrm{V\,m^{-1}}$, respectively, but still persisted. Concretely, the systematic effects taken into account were a substantial modulation in PMT detection efficiency due to the trap RF drive (see §5.6.1) and an additional modulation from magnetic fields at the ion position presumably combined with at least one other effect (see §5.6.3). With the exception of a $\sim 1\,\mathrm{V\,m^{-1}}$ shift in some cases far outside the linear regime, no evidence suggesting any systematic shifts affecting the parametric excitation method was found (§5.6.4).

Ultimately, the origin of these remaining systematic shifts remains unclear; a discrepancy between RF correlation and sideband methods of similar magnitude was reported in ref. [Kel+15]. As the parametric excitation method probes the position to the RF null very directly through modulation of the RF potential, yet does not rely on demodulation of any measurements near that frequency, it should be fairly robust against many potential systematics. Even a residual offset on the few-$\mathrm{V\,m^{-1}}$ level would not be an issue for our use case (see §4.5.1). To gain certainty about the origin of the remaining shifts, and validate the parametric modulation method for use in more sensitive applications, it would be still useful to implement another high-resolution method in our apparatus for comparison, such as that recently proposed in ref. [Hig+21].

5.6.1   RF correlations: Photomultiplier gain modulation

As previously reported by Keller et al. [Kel+15], we find that the quantum efficiency of the photomultiplier tube used to detect the ion fluorescence is modulated by the presence of ambient radio-frequency fields, in particular by radiation leaking from the trap electrodes. Extraneous frequency content near the carrier frequency has the potential to cause systematic shifts in any measurement relying on synchronous demodulation, such as the RF correlation and parametric excitation micromotion compensation schemes:

If the detection probability of photons is modulated by $G(t) = 1 + \mathrm{re}(g\, e^{i\Omega_{rf}t})$, where g is the complex modulation amplitude, a signal of amplitude $r \in \mathbb{C}$ at frequency Ω_{rf} in the fluorescence rate $R(t) = R_0(1 + \mathrm{re}(r\, e^{i\Omega_{rf}t}))$ will yield a demodulated amplitude of $\langle G(t)R(t)e^{-i\Omega_{rf}t}\rangle_t \propto r + g$ instead, i.e. the efficiency modulation results in a shift of the measured signal in the complex plane. In the context of stray field compensation, a component of g along r will cause a systematic offset in the inferred minimum.

The phase difference between r and g depends on various phase shifts and time delays, whether in the RF signal chain (e.g. the physical detector placement) or the ion phase response (influenced by the cooling laser parameters). Thus, there is no reason to believe that the gain modulation shift should be orthogonal to the signal direction, and hence must be characterised so it can be subtracted from the data in order to obtain accurate results.

For the data presented here, we use a Hamamatsu *H10682-210* photomultiplier module to detect ion fluorescence, situated approximately 0.5 m from the trap centre. With the ion trap empty, light from a torch is leaked into the apparatus to provide photon counts at a rate close to that used for the RF correlation measurements, but certainly free of correlations with the trap RF drive. A modulation with a magnitude of 1.2 % is measured; it appears to be fractionally constant across a range of signal count rates, and is not present in the (significantly lower) dark count rate. Subtracting this offset from from the RF correlation data leads to a shift of 2.3 V m^{-1} in the estimated micromotion minimum. As discussed at the beginning of this section, a small difference between the results of the RF correlation and parametric excitation methods remains even taking this offset into account.

While the exact origin of this effect is not known, it is plausible that the parametric excitation method proposed here should be far less susceptible to it: the RF amplitude is only modulated by a small amount $h \ll 1$ %, and as long as the detection chain responds linearly, the signal will only exhibit frequency components at $\Omega_{rf} \pm \omega_{am}$ (but not at ω_{am}, where it is demodulated to extract the motional amplitude). Even for a large modulation depth $h = 6$ %, the correlated amplitude for background light was indeed measured to be consistent with zero, $S = 6(11) \times 10^{-5}$.

5.6.2   *Radiation pressure*

Radiation pressure from 422 nm photons can be excluded as a reason for the mismatch in our parameter regime. For a typical count rate of 15 kHz observed in the parametric excitation and RF correlation measurements and a detection efficiency of 3.6×10^{-4}, the ion scatters $\gamma = 4.0 \times 10^6$ s^{-1} photons. Since the photons are re-emitted isotropically, the average force experienced by the ion is simply

$$\boldsymbol{F} = \hbar\,\gamma\,\boldsymbol{k}.$$

For the above parameters, $|\boldsymbol{F}| = 6.3$ zN, which equivalent to an electric field of $|\boldsymbol{F}|/e = 39$ mV m^{-1}. This estimate does not include the projection onto the compensation direction
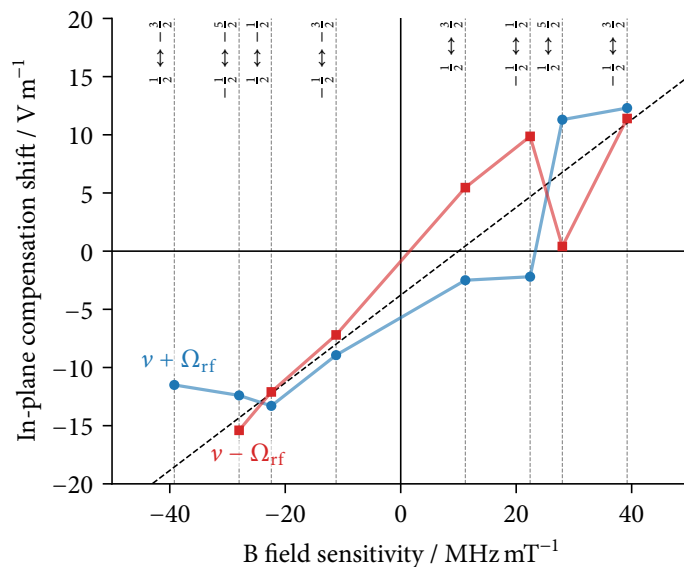
Figure 5.12: In-plane compensation fields that minimise the observed Rabi frequency on the $\pm 1^{\text{st}}$ micro-motion sidebands of each of the 674 nm electric quadrupole transitions (connecting $S_{1/2}$ and $D_{5/2}$ Zeeman sublevels) accessible in the experiment. The transitions are arranged by their sensitivity to magnetic fields; the labels on the top denote the respective m_J quantum numbers. The measurement uncertainty for each point is smaller than the marker size; the points are connected by lines to guide the eye. The dashed line shows a least-squares fit to the mean of the two measurements for each transition. The fields are given relative to the compensation field obtained from the parametric excitation method, which was stable to $< 1 \, \text{V m}^{-1}$ over the measurement duration.

of interest, which attenuates the effect by a factor of ≈ 0.7 in this comparison. While our method is capable of resolving this effect (see §5.5), it is far too small to explain the observed discrepancies.

5.6.3 674 nm sideband Rabi frequency: transition dependence

As the micromotion sideband suppression method works by minimising the lab-frame modulation of a narrow-linewidth transition by micromotion at the trap frequency, it is affected by other mechanisms that modulate the transition frequency at Ω_{rf}.[10] The 674 nm $S_{1/2} \leftrightarrow D_{5/2}$ electric quadrupole transitions we use for this are sensitive to magnetic fields through the Zeeman effect. One possible modulation source are thus residual magnetic fields at the ion position caused by the trap RF electrode supply paths, as recently described in ref. [Mei+18]. This effect would scale with the strength of the oscillating magnetic field at the ion and the Zeeman sensitivity of the probed transition.

When comparing the points of minimum excitation across the different 674 nm transitions accessible in our experiment, as shown in fig. 5.12, we indeed observe a systematic shift roughly consistent with a linear dependence on the magnetic field sensitivity. We,

---

[10]This is much the same as for the RF correlation and parametric excitation methods, only that the signal is "demodulated" here by the ion, which effectively probes the field in its rest frame.

however, also observe at least one additional effect that leads to different shifts depending on the sign of the micromotion sideband probed, with the effect varying across the different transitions in a complex fashion.

The compensation value quoted in the above comparison to the other methods was obtained as the interpolation to zero from a linear fit to the data from fig. 5.12, that is, effectively averaged across Zeeman transitions and sideband orders. However, until the additional mechanism(s) at play can be elucidated through a more systematic investigation[11], and the symmetry of the resulting shifts confirmed, the accuracy of the results from the 674 nm sideband method quoted here should be considered to be limited to $\sim 5 \, \text{V m}^{-1}$.

5.6.4   Parametric excitation method

Benign shift from out-of-phase RF field. The dataset from fig. 5.7 shows an effect not captured by the pseudopotential approximation: Instead of crossing the origin, the line describing the correlated amplitude is offset from zero by a small amount, here corresponding to a correlated fraction of 0.006 (a displacement equivalent to a stray field of $0.12 \, \text{V m}^{-1}$ along the fit line). Numerical simulations of eq. (5.1), shown in fig. 5.13, with an added position-independent $\sin(\Omega_{\text{rf}} t)$ term indicate that a possible origin would lie in a residual out-of-phase RF field of $\sim 60 \, \text{V m}^{-1}$ near the ion position, for instance caused by slight imbalances in drive or grounding of the two RF electrodes. Similar to the equivalent effect in RF correlation measurements, for which an analytical approximation is more readily derived [Ber+98; Kel+15], this appears to shift the observed complex rate in a direction precisely perpendicular to the modulation caused by stray DC fields, thus not affecting the inferred compensation optimum.

Partial phase modulation from non-ideal modulation circuitry. Both from path-length mismatches in trap RF modulation circuitry (see §3.2.1) and from a mismatch in the resonant impedance matching circuit response between the sidebands (e.g. if Ω_{rf} is not perfectly centred on its resonance), it would be possible for the trap RF drive to also acquire some phase modulation in addition to the desired amplitude modulation component. Figure 5.14 shows a numerical simulation of the full equations of motion from eq. (5.1) for a range of static offset fields, showing that the response changes phase and loses in amplitude if some of the modulation is distributed into the phase channel, but does not shift the position of the signal null. As such, this would not affect the stray field calibration accuracy, although we typically operate our apparatus in a way producing very close to pure amplitude modulation anyway.

---

[11] It is well possible that this is one of a number of effects recently discussed in the optical clock literature [Gan+18; Arn+19].
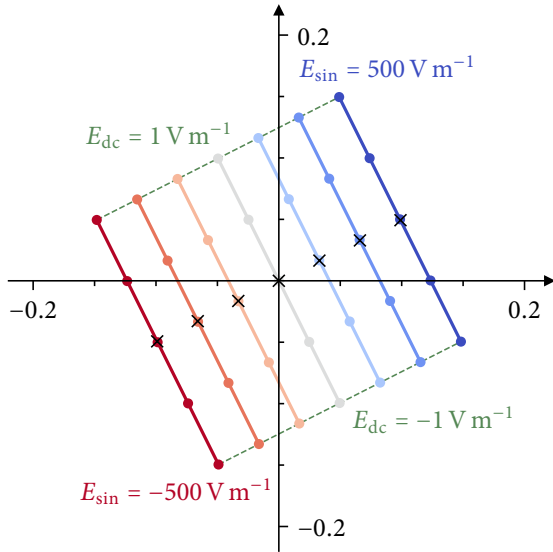
Figure 5.13: Response of the parametric excitation correlation signal to the presence of a uniform out-of-phase RF term $E_{sin} \sin(\Omega_{rf}t)$ in addition to the $\cos(\Omega_{rf}t)$ quadrupole trap potential. Each line shows the complex amplitude of the ion's velocity at ω_{am} at one choice of E_{sin} as the stray field is varied, with the markers being equally spaced between offset fields $E_{dc} = -1\,\mathrm{V\,m^{-1}}$ and $+1\,\mathrm{V\,m^{-1}}$. The extra field results in a linear shift in the demodulated signal, but does not affect stray field compensation accuracy, as it is perpendicular in direction.
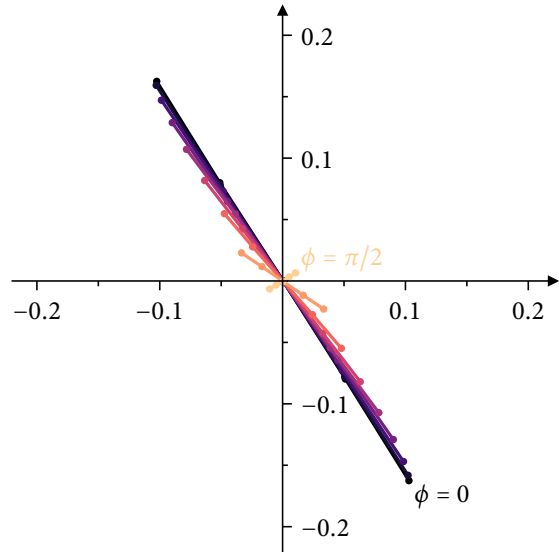
Figure 5.14: Response of the parametric excitation correlation signal to a phase shift ϕ between constant trap RF drive $\cos(\Omega_{rf}t)$ and modulation term $h\cos(\omega_{am}t)\cos(\Omega_{rf}t + \varphi)$, for $h = 0.01$. Each line shows the complex amplitude of the ion's velocity at ω_{am} at one choice of phase offset φ as the stray field is varied, with the markers being equally spaced between offset fields $E_{DC} = -1\,\mathrm{V\,m^{-1}}$ and $+1\,\mathrm{V\,m^{-1}}$. The phase offsets range from pure amplitude modulation at $\varphi = 0$ to maximally mixed modulation at $\varphi = \pi/2$, where the correlated signal is greatly diminished in amplitude, but the stray-field-less case is still in the origin.

Frequency-dependent shift for large modulation amplitudes. According to the theoretical treatment developed in §5.2, the stray field compensation optimum is unambiguously given by the point where the correlated amplitude vanishes. The modulation frequency only affects the strength (and, close to resonance, phase) of the correlation signal.

At low and moderate modulation depths h, the correlation optimum, as inferred by the complex linear regression procedure, is indeed observed to be independent of the choice of modulation frequency ω_{am}, enabling robust stray field compensation. At high modulation depths, however, we observe a systematic shift of up to $1\,\mathrm{V\,m^{-1}}$ (compared to the value obtained at lower excitation strengths) for ω_{am} slightly above the mode frequency. An example of this behaviour is shown in fig. 5.15.

The exact mechanism for this is unknown and the behaviour could not be replicated in numerical simulations based on eq. (5.1). Empirically, however, we only observe this shift if the ion is driven very strongly to where the response is no longer linear and, for ω_{am} close to the mode frequency, a parametric instability is excited. This can be diagnosed by its effect on the total fluorescence rate, which remains constant in the linear regime.
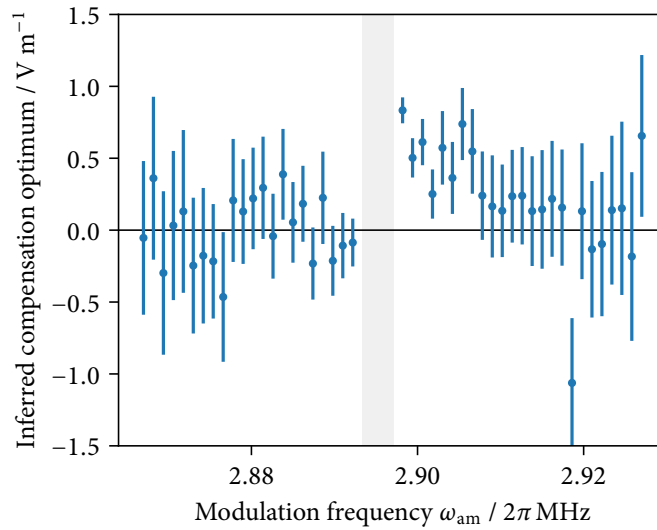
Figure 5.15: Systematic shift in the inferred compensation optimum at large excitation amplitudes: Each point represents the compensation optimum inferred from a fit to 51 correlation measurements at offsets of $\pm 10\,\mathrm{V\,m^{-1}}$ along the HF mode direction, with modulation index $h \approx 0.02$ and cooling parameters $\Delta = -2\pi \times 13\,\mathrm{MHz}$, $s \approx 0.9$. For modulation immediately above the mode frequency (the shaded region includes measurement uncertainty and drifts over the measurement duration), a systematic shift towards positive values is apparent.

By choosing the experimental parameters such that the fluorescence does not significantly change as ω_{am} is scanned over the resonance, or as the modulation is turned on and off at the target ω_{am}, the problematic parameter regime can easily be avoided in practice.

5.7   Discussion

We have described a method for stray field compensation in radio-frequency traps based on parametric excitation of the secular modes of motion and synchronous detection of particle motion. No directional resolution is required for the motional measurement, as directions are addressed via spectrally separated oscillatory modes. This is particularly convenient for surface-electrode traps, for which optical access may be limited. For trapped ions, this means a single laser cooling beam is sufficient for compensation in all spatial directions as long as it has a projection onto all modes, which is in any case required for efficient laser cooling. Neither the ability to access transitions of linewidths much smaller than the trap drive frequency nor to perform cooling to the motional ground state is necessary. The presented method is thus equally applicable to other types of RF-trapped systems, as long as motion can be weakly damped and detected.

Phase sensitivity in the detection of motion is not a strict requirement; the compensated point could be found by just minimising the magnitude of parametrically excited motion. When phase information is available, however, the compensation optimum can

trivially be extracted in a robust fashion even in the presence of large amounts of statistical noise by linear regression, making this technique well-suited to unattended calibration routines. Compared to approaches based on variations in the total fluorescence rate due to parametric instabilities, where the change in rate away from the RF null can be positive or negative depending on the parameters, the signal here is unambiguous and does not require careful development of robust fitting algorithms. As the method utilises excitation amplitudes below the parametric instability threshold and works in a small-amplitude regime where the cooling dynamics are to good approximation linear, the trap stability is hardly affected; we do not observe elevated ion loss rates during stray field calibrations.

The compensation method is simple to implement in practice. Technical requirements are limited to a provision for amplitude modulation of the trap potential, whether through a built-in function of the trap RF source or an external circuit. The required modulation depth is small, so the necessary amplitude-modulation sidebands can be passed through all but the most narrow-band step-up resonator designs (or replaced by a baseband signal directly to the RF electrodes). Compared to techniques which detect stray fields by imaging the shift in ion position resulting from changes in RF trap strengths, an advantage of this scheme is that the trap configuration can be kept nearly constant, alleviating any concerns about systematic shifts due to changes e.g. in RF power dissipation. The required photon timestamp resolution is set by the motional frequencies; with typical secular mode frequencies in the low MHz region, phase jitter small enough that it does not significantly reduce the signal can easily be achieved using almost all commercially available photon detectors and real-time control hardware.

While the method relies on knowledge of well-separated secular mode frequencies, the width of the parametric resonance feature can be tuned via the cooling parameters. For our parameters, typical linewidths are on order of 1 kHz, which is workable, even though our apparatus has not been designed with a focus on radial mode stability or a particularly harmonic trapping potential, and we indeed observe mode frequency changes of several kilohertz due to laboratory temperature fluctuations.

The achieved statistical uncertainties are comparable to or below that of the lowest results reported in the literature for other stray field compensation methods (see table 5.1). In addition to finding the stray field minimum, this method can in principle also be used to directly obtain single-point estimates of the RF potential at the ion position owing to the linear response. For this application, however, either a detailed description of the ion cooling dynamics or the capability to apply well-characterised additional fields for calibration (i.e. a locally accurate model for the DC electrode fields) is necessary to calculate the scaling factor.

While numerical results indicate a small sensitivity to out-of-phase RF field terms such as caused by phase asymmetries in the RF electrode drive, this cannot be described in the simple pseudopotential picture presented here. This method is hence not particularly suit-

	σ_E $\mathrm{Vm^{-1}}$	α $\mathrm{Vm^{-1}/\sqrt{Hz}}$	OOP comp. in our setup
Cavity emission spectroscopy [Chu+13]	1.8	1.6	$\times$
Repumper RF correlations [All+10]	1	1	$\times$
Micromotion sideband Rabi frequency [Chw09]	0.4	-	$\times$
Parametric excitation (total fluorescence rate) [Kel+15]	0.3	-	$\checkmark$
Focus-scanning position shift imaging [Glo+15]	0.09	2.30	$\sim$
RF correlations [Kel+15]	0.09	0.57	$\times$
Co-trapped ultracold atoms [Här+13]	0.02	-	$\times$
Direct position shift imaging [Hir+20]	≈ 0.01	-	$\times$
This method	0.015	0.1	$\checkmark$
Optical dipole trap beam overlap [Hub+14]	0.009	≈ 0.08	$\times$
Ramsey measurement of position shift [Hig+21]	0.0035	0.031	$\times$

Table 5.1: Comparison of stray field uncertainties achieved using various compensation methods, and whether they would be suitable for addressing the out-of-plane direction in our apparatus (without major hardware additions). To provide an approximate indication of the performance achieved in typical experiments, the lowest published uncertainties σ_E for each method are given, along with the inverse sensitivity α (where reported) for the same measurement. Note that the figures heavily depend on ion species and experimental parameters, as do the conversion factors to other quantities of interest (e.g. the induced second-order Doppler shift); see ref. [Kel+15] for a more detailed comparison of some of these results.

able for the investigation of "quadrature" excess micromotion caused by such effects, the compensation of which in either case requires changes to the RF trapping fields [Ber+98; Moh+19].

Here, we have only discussed a mode of operation where both parametric excitation close to the motional frequency ($\omega_{\mathrm{am}} \approx \omega_0$) and cooling are applied in a stationary fashion, as the resulting signal is robust, linear in the stray field, and independent of initial conditions. To increase sensitivity if the available modulation depth is limited, this scheme could be combined with excitation at the principal parametric resonance ($\omega_{\mathrm{am}} \approx 2\omega_0$) to amplify small amounts of motion [YDN93]. By preparing the motional modes in a sufficiently well-known state through initial coherent excitation (possibly after ground state cooling), parametric excitation could also be used subsequently to probe stray fields where no dissipation mechanism is available, for instance along laser-free storage or interaction regions of an extended shuttling-based quantum computing chip. It should also be possible to acquire data for multiple mode directions in parallel by applying multiple amplitude-modulation tones at the same time. The method would scale to multiple-ion crystals via use of their centre-of-mass (in-phase) motional modes.

In conclusion, we have presented a simple method for multi-directional stray field compensation in Paul traps based on spectrally addressed parametric excitation of the secular modes through amplitude modulation of the RF trapping potential. The technique is applicable to a wide variety of trapped particle systems; in typical trapped-ion experiments, the added technical complexity is minimal, and the achieved resolution compares favourably with that reported for alternative methods.

6 Experimental entanglement characterisation

We now turn to the characterisation of the entanglement generation primitives – between a single ion and a single photon, and after entanglement swapping, between two remote ions – in the experiment. We begin with a look at the correlations between the ion state after spontaneous emission and the polarisation of emitted photons (§6.1), where we will see two distinctive geometric effects predicted by the theory from §4.2: the shift in apparent emitter position for σ-polarised photons, and the avoidance of polarisation mixing if the single-mode fibre is positioned exactly on the optical axis. §6.2 then describes the joint quantum state between ion and photon polarisation including coherences through the use of state tomography. §6.3 finally presents results from the characterisation of the entangled state between two remote ions, again through state tomography.

Our perspective throughout will be one of evaluating the remote entanglement link as a building block for more complex quantum networking applications. As such, we at no point attempt to back out any known imperfections from the experimental data. The latter is sometimes done in studies of individual mechanisms (e.g. "readout normalisation" in work on high-fidelity gates), but is of course impossible in many applications where the entanglement is further used in some way. As such, the presented data place a lower bound on the fidelity achievable through this remote entanglement generation mechanism. High-fidelity local operations are not difficult to achieve in trapped ions, however, and as discussed at the end of the chapter in §6.4, our results significantly improve on remote entanglement quality across all matter qubit platforms.

6.1 Ion–photon correlations

Evaluating the ion–photon polarisation correlations is somewhat challenging for two fundamental reasons: first, it not possible to obtain spontaneous emission from a $J = 1/2 \leftrightarrow J' = 1/2$ transition that is purely π or σ polarised, and secondly, the optical fibre necessary as a spatial filter of the unwanted non-orthogonal polarisation components will always introduce an uncontrolled polarisation rotation due to birefringence[1]. To establish a po-

---

[1] One could attempt to manually compensate it by coupling cw laser light through the system to align the polarisation measurement bases with the magnetic field direction. This would be cumbersome to achieve with a high degree of precision, though.
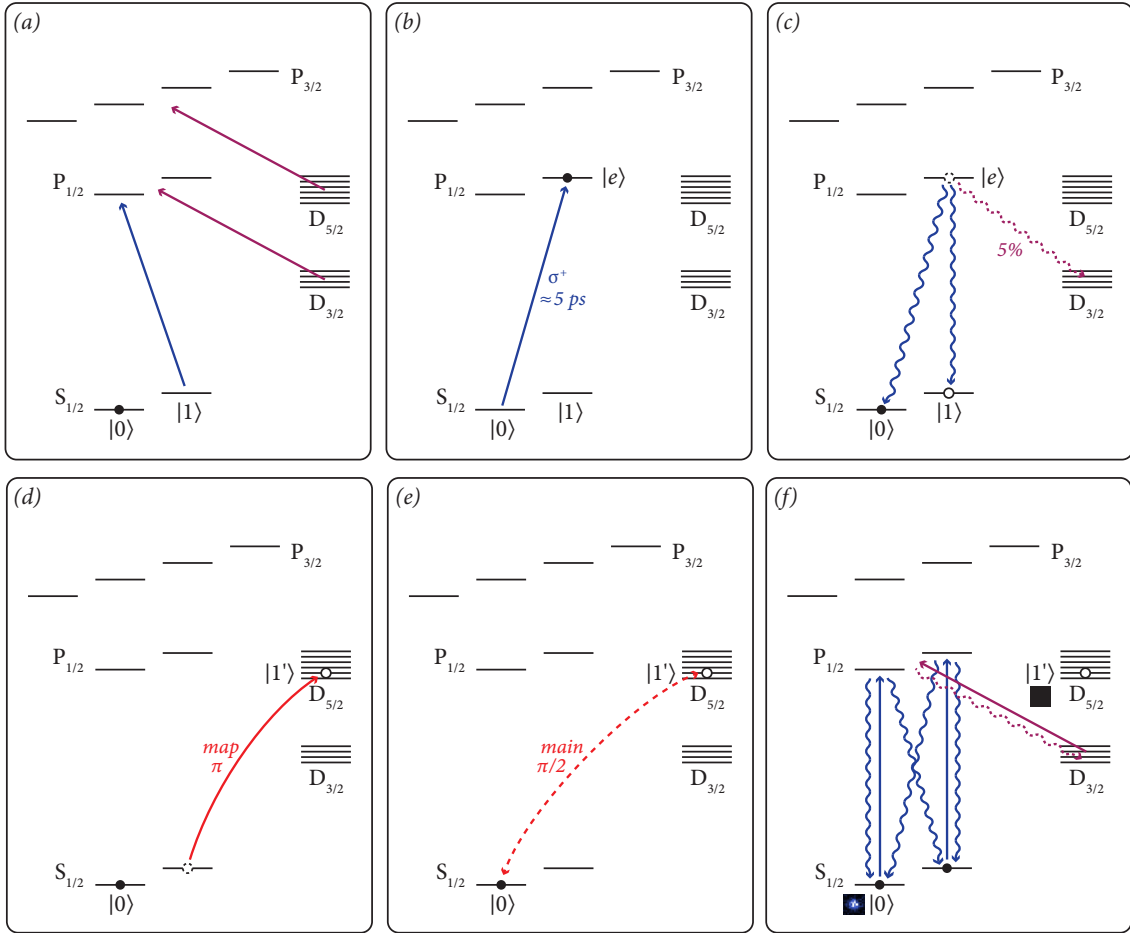
Figure 6.1: Basic steps in the experiments for ^{88}Sr$^+$ ion–photon polarisation correlation analysis, ion–photon state tomography, and the ion–ion state tomography experiments (for which this is executed by both nodes with shared heralding logic). (a) The ion is initialised in $|0\rangle \coloneqq \left|5s\,^2S_{1/2}, m_J = -\frac{1}{2}\right\rangle$ through optical pumping; infrared lasers repump any population trapped in the metastable D levels. (b) A resonant ~ 5 ps laser pulse excites the ion to $|e\rangle \coloneqq \left|5p\,^2P_{1/2}, m_J = \frac{1}{2}\right\rangle$. (c) The ion spontaneously decays from $|e\rangle$, with high probability into one of the Zeeman-split ground states. (d) A π pulse maps $|1\rangle \coloneqq \left|5s\,^2S_{1/2}, m_J = \frac{1}{2}\right\rangle$ to $|1'\rangle \coloneqq \left|4d\,^2D_{5/2}, m_J = -\frac{3}{2}\right\rangle$. (e) If the ion state is to be read out along an axis in the xy plane of the Bloch sphere instead of the computational basis, a $\pi/2$ pulse on the "main" optical transition between $|0\rangle$ and $|1'\rangle$ is applied. (f) The ion state is read out using resonant fluorescence detection.

larisation measurement aligned with the ion decay channels, one thus needs to work in the single-photon regime and measure the ion state after each photon detection to determine which of the two transitions the photon originated from.

6.1.1   Experimental sequence

The steps in the experimental data acquisition cycle for such experiments with single spontaneous decay photons are outlined in fig. 6.1. First, the ion is prepared in the $|0\rangle \coloneqq \left|5s\,^2S_{1/2}, m_J = -\frac{1}{2}\right\rangle$ ground state through optical pumping with σ^--polarised 422 nm light. It is then excited to $|e\rangle \coloneqq \left|5p\,^2P_{1/2}, m_J = \frac{1}{2}\right\rangle$ by a σ^+-polarised 422 nm ~ 5 ps laser pulse.
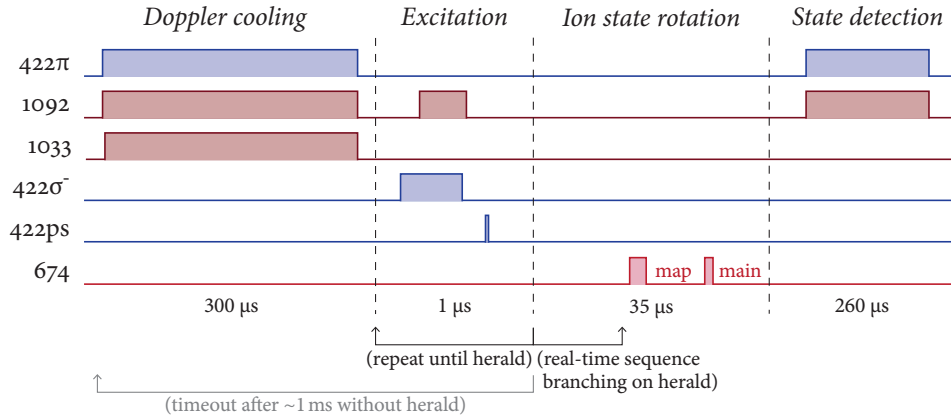
Figure 6.2: Timing of the experimental sequence (not to scale) for ion–photon polarisation correlation analysis, ion–photon state tomography, and ion–ion state tomography experiments. The first 674 nm laser pulse on the "map" transition $\left|5s\,^2S_{1/2}, m_J = \frac{1}{2}\right\rangle \rightarrow \left|4d\,^2D_{5/2}, m_J = -\frac{3}{2}\right\rangle$ is always present. For measurement away from the computational basis, an additional pulse on the "main" transition $\left|5s\,^2S_{1/2}, m_J = -\frac{1}{2}\right\rangle \leftrightarrow \left|4d\,^2D_{5/2}, m_J = -\frac{3}{2}\right\rangle$ further rotates the ion state before state-dependent fluorescence readout. The sequence of state manipulation and detection pulses to run can be selected in real time depending on the detector click pattern observed at the herald station (which is not necessary for state tomography, but is relied on in the DIQKD experiments described in chapter 7).

This excitation sequence is coordinated by the real-time FPGA state machine, such that it can be retried with minimal delays if no 422 nm spontaneous emission photon is detected (attempt rate: 1 MHz). If a photon is in fact registered on the detectors – we use the same Bell-state measurement apparatus throughout, also for single-ion experiments –, we map $|1\rangle := \left|5s\,^2S_{1/2}, m_J = \frac{1}{2}\right\rangle$ to the $|1'\rangle := \left|4d\,^2D_{5/2}, m_J = -\frac{3}{2}\right\rangle$ state using a 674 nm laser pulse. Like this, the resulting state can be read out using fluorescence detection on the $S_{1/2} \leftrightarrow P_{1/2} \leftrightarrow D_{3/2}$ cycle and recorded, after which the ion is Doppler-cooled again, and the next iteration begins. A Doppler-cooling period is also inserted if no herald is detected in a certain period of time to avoid excessive ion heating (typically 1 ms), although this limit is seldom reached in typical single-ion experiments on account of the total system efficiencies.

This sequence structure is also shown as a timing diagram in fig. 6.2. This basic structure remains the same for all single-ion experiments described in this chapter, with the ion qubit manipulation 674 nm laser pulses chosen as required to implement state readout in the desired basis. It is also the sequence run on each node independently and in parallel for the two-node ion–ion entanglement experiments described here and in chapter 7, for which the FPGA state machines on each node coordinate between each other to synchronise timing and heralding success conditions.

6.1.2   *Aligning the polarisation measurement basis*

In the Bell-state measurement apparatus used also for all the single-ion experiments, the photonic polarisation measurement basis set by the PBS cubes is rotated by (in addition to any residual birefringence in the optics) the two wave plates following the optical fibres from the trap nodes. As discussed in §2.3.2, they effect a rotation $U_{\mathrm{wp}}(\alpha_2, \beta_2) \cdot U_{\mathrm{wp}}(\alpha_4, \beta_4)$ on the incoming photon polarisation, where $\alpha_i \in [0, 2\pi)$ are the retardances, $\beta_i \in [0, \pi)$ the orientations of the fast axis w.r.t. the horizontal plane, and $i = 2$ denotes the half-wave plate, $i = 4$ the quarter-wave plate. A convenient parametrisation[2] of the angles is through the parameters $\{\gamma_1, \gamma_2\} \subset \mathbb{R}$ via $\beta_2 := \frac{\gamma_1 + \gamma_2}{4}$, $\beta_4 := \frac{\gamma_1}{2}$. Then, for ideal wave plates ($\alpha_2 = \pi$, $\alpha_4 = \frac{\pi}{2}$), the combined unitary is

$$
\begin{aligned}
U &:= U_{\mathrm{wp}}\left(\pi, \frac{\gamma_1 + \gamma_2}{4}\right) \cdot U_{\mathrm{wp}}\left(\frac{\pi}{2}, \frac{\gamma_1}{2}\right) \\
&= \begin{pmatrix} \cos\gamma_1 \cos\gamma_2 + i\sin\gamma_1 \sin\gamma_2 & \sin\gamma_1 \cos\gamma_2 - i\cos\gamma_1 \sin\gamma_2 \\ \cos\gamma_1 \sin\gamma_2 - i\sin\gamma_1 \cos\gamma_2 & \sin\gamma_1 \sin\gamma_2 + i\cos\gamma_1 \cos\gamma_2 \end{pmatrix}.
\end{aligned}
\tag{6.1}
$$

The measurement corresponding to PBS transmission represented as a Bloch vector $\xi(\rho)$ is then

$$
\xi\left(U^\dagger \, |0\rangle\langle 0| \, U\right) = \begin{pmatrix} \sin(2\gamma_1)\cos(2\gamma_2) \\ \sin(2\gamma_2) \\ \cos(2\gamma_1)\cos(2\gamma_2) \end{pmatrix},
\tag{6.2}
$$

i.e. $2\gamma_1$ and $2\gamma_2$ are spherical coordinates, where the polar axis contains the circular polarisation states and the origin of the azimuthal coordinate γ_2 is chosen such that $\gamma_2 = 0$ describes the equator, not one of the poles. In the experiment, the retardance of the wave plates will differ slightly from the ideal values, so the correspondence to spherical coordinates will only be approximate. We nevertheless parametrise the angles through γ_1 and γ_2 for interactive exploration/iterative optimisation in the experiment, as this still aids ease of interpretation as long as the wave plates are reasonably accurate (and compute the accurate unitaries used for ion–photon state tomography from these parameters from the actual retardances as determined using a crossed-polariser measurement).

The polarisations corresponding to the atomic decay channels, as rotated by the fibre birefringence, can simply be found by empirically optimising γ_1 and γ_2 such that the ion state post-selected on detection of a photon on a given APD is as close to the computational basis states as possible. This proceeds as described in §6.1.1, without any additional qubit manipulation pulses apart from the transfer of the $|1\rangle$ Zeeman state to the $D_{5/2}$ manifold for state readout; a sample scan is shown in fig. 6.3.

The assignment between detector channels and atomic decay channels can be chosen arbitrarily as long as it is the same between the ion trap nodes. We typically choose to align the σ^+ decays, which remain bright in the state readout, with transmission on the PBSes, as

---

[2]This parametrisation was originally suggested by P. Drmota based on empirical observations.
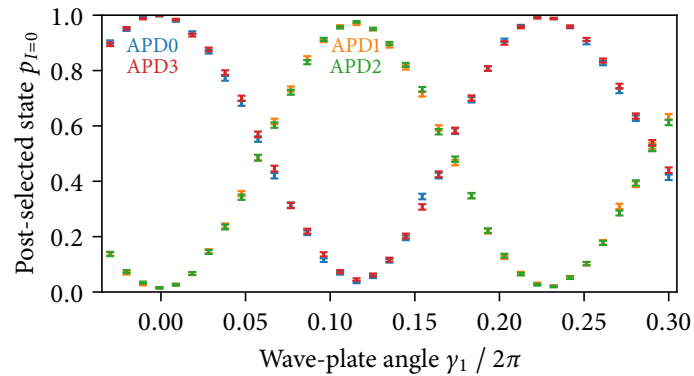
Figure 6.3: Ion state after detection of a photon in the given detector for a varying wave plate rotation γ_1 (see text for definition); 10 000 photons total for each angle acquired in node Alice. The symbols show the probability to find the ion in the non-shelved state $|0\rangle$ after detection of a photon in the given APD (with the 1σ confidence interval from binomial statistics). The reduced contrast in the middle extrema is caused by the fact that the retardance of the wave plates in the apparatus differs slightly from the ideal values.
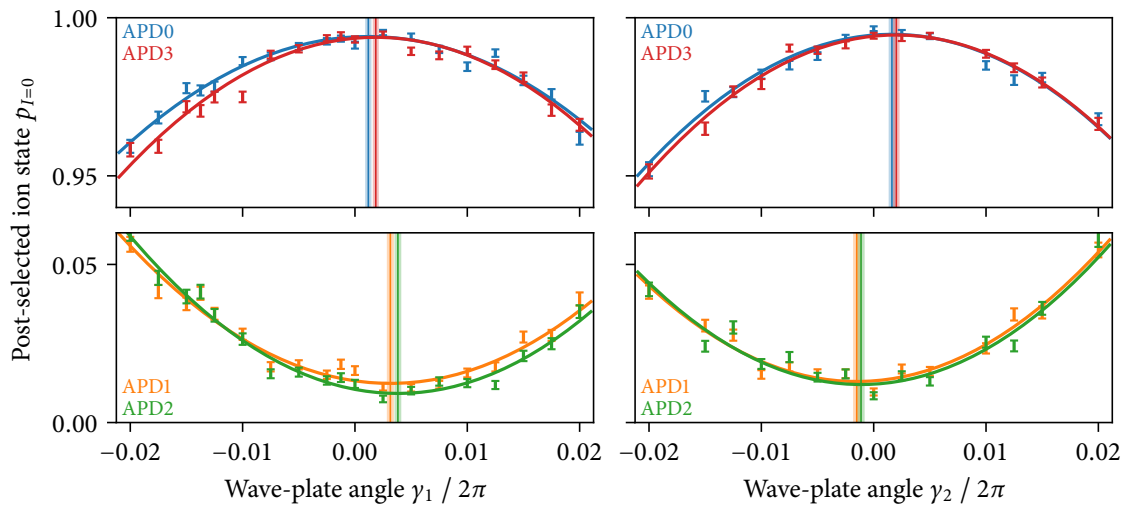
Figure 6.4: Ion state after detection of a photon in the given detector for varying wave plate rotation angles γ_1 and γ_2 (see text) close to a previously determined optimum; 10 000 photons total for each angle acquired in node Bob. The symbols show the probability to find the ion in the non-shelved state $|0\rangle$ after detection of a photon in the given APD (with the 1σ confidence interval from binomial statistics). Parabolic fits are shown as solid lines, with the vertical lines denoting the position of the extremum and its standard error from the least-squares fit. A statistically significant offset is visible between the extrema in the results post-selected on PBS transmission and reflection; at the time of writing, its origin was unclear.
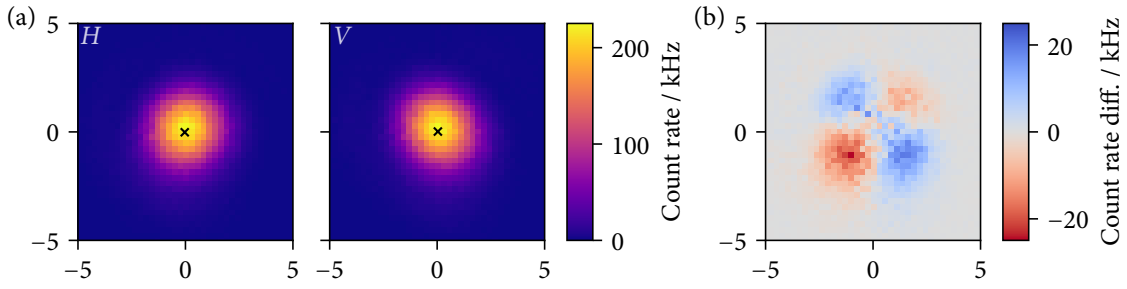
Figure 6.5: Spatial dependence of steady-state fluorescence count rates. The axes give the collection fibre tip position in micrometres, with the direction matching the laboratory orientation shown in fig. 3.6 (the ambient magnetic field points from top left to bottom right). *(a):* Total count rates across the two APDs corresponding to each polarisation (set up such that $\sigma^{\pm} \to H, \pi \to V$ on axis). Crosses mark the centre of mass of the distribution, which is virtually identical between the polarisations, although their shape differs slightly. *(b):* Difference between the count rates, emphasising their different spatial dependence. Blue marks the regions where V ($\parallel \boldsymbol{B}$) polarisation dominates, red those where H ($\perp \boldsymbol{B}$) does.

this gives the signal of the highest contrast if some $\sigma \to \pi$ polarisation mixing, finite PBS extinction ratios, or imperfections in the 674 nm laser pulses or state readout are present (though these effects are small). A narrow scan close to an optimum is shown in fig. 6.4. Curiously, we consistently observe a small but statistically significant offset between the optima obtained from the detectors in PBS transmission or PBS reflection; at the time of writing, its origin remained unclear.

6.1.3 *Position-dependent count rates for steady-state fluorescence*

With the wave plate set such that the detectors are aligned with the atomic polarisation axes, we can investigate polarisation effects in the position dependence of the steady-state fluorescence of the ion during Doppler cooling on the $S_{1/2} \leftrightarrow P_{1/2} \leftrightarrow D_{3/2}$ cycle, which we also use to optimise the fibre position to match the ion image. A sample scan in node Bob is shown in fig. 6.5. The two σ decays occur with equal rate; this is necessarily the case, as the 422 nm and 1092 nm laser beams, incident perpendicular to the magnetic field, do not carry any angular momentum. As expected from the Clebsch–Gordan coefficients and dipole emission patterns (§4.1), H photons (set up to correspond to σ) and V photons (set up to correspond to π) are both detected with approximately equal rate. The shapes differ, however: As the H detection channel captures photons scattered on both σ^+ and σ^- transitions, it is elongated in the direction perpendicular to the magnetic field as a consequence of the apparent positional shift inherent in observing circularly polarised dipole emitters described in §4.2.3 (σ^+ and σ^- are each shifted slightly to opposite sides, leading their incoherent sum to take an elongated shape). The V channel peak is slightly elongated in the direction along the magnetic field, as it captures contributions from the longitudinal components of the σ dipoles there. As the centres of mass of the distributions (or the centres
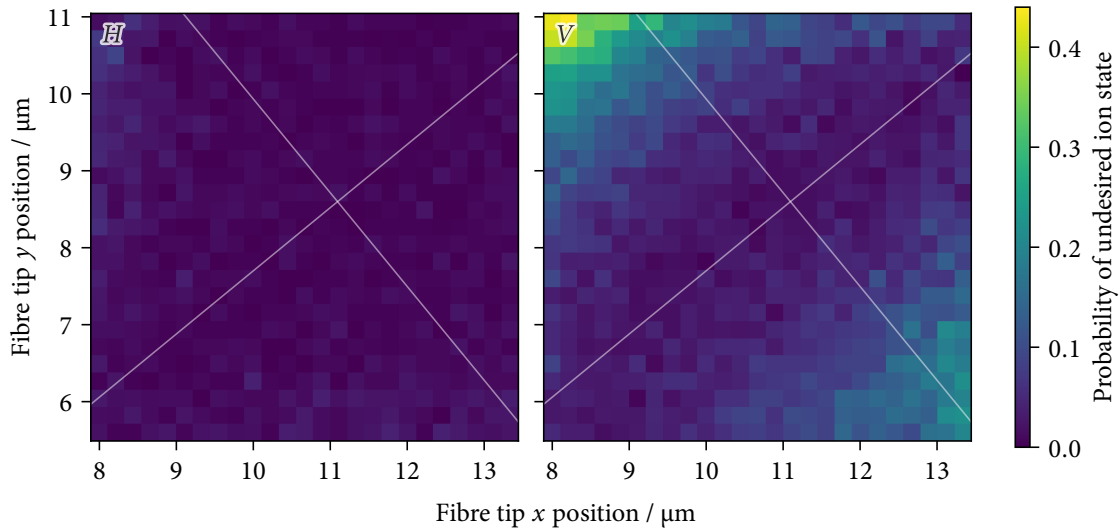
Figure 6.6: Ion state following detection of a photon of H and V polarisation (set up such that $\sigma^\pm \rightarrow H$, $\pi \rightarrow V$ on axis), given as the probability to obtain the undesired outcome during the fluorescence measurement at the end of the sequence (ideally, $H \leftrightarrow |0\rangle$ (ion bright), $V \leftrightarrow |1\rangle$ (ion dark) with unit probability). For each of the 25×25 points, scanned in lines from bottom to top, 500 clicks were collected between all four APDs. The x and y axes match the laboratory orientation shown in fig. 3.6 with the ambient magnetic field pointing diagonally in the down-right direction. The white crosses mark the directions parallel and perpendicular to the magnetic field (inferred from a parabolic fit to the V channel) centred on the optical axis (inferred from the π collection efficiency maximum, see fig. 6.8).

of a two-dimensional Gaussian fit) still match, however, special care is not required when optimising the fibre position in the experiment.

6.1.4 Position-dependent ion–photon correlations

To study these effects depending on the position of the collection fibre in more detail, we return to the single-photon setting, where the signal measured is the ion state after detection of a spontaneously emitted photon following picosecond laser excitation.

The dependence of the ion state following detection of a photon on a detector calibrated to correspond to σ or π polarisation on the transversal position of the fibre tip is shown in fig. 6.6. As expected from theory (see §4.3.1), the cancellation of the longitudinal dipole contribution that causes $\sigma \rightarrow \pi$ polarisation mixing only works in the plane perpendicular to the magnetic field. Displacement of the fibre in either direction along the magnetic field causes the ion to be observed in $|0\rangle$ even following the detection of a photon on an APD set up to correspond to π polarisation. A displacement perpendicular to the magnetic field does not appreciably change the post-selected ion states, though. This is shown more explicitly in fig. 6.7 along with the theory prediction.

By combining the post-selected ion state with the number of attempts until detection of a photon and relative frequency of the four detector clicks in these measurements, it is also possible to map out the overall collection efficiency for photons from either atomic
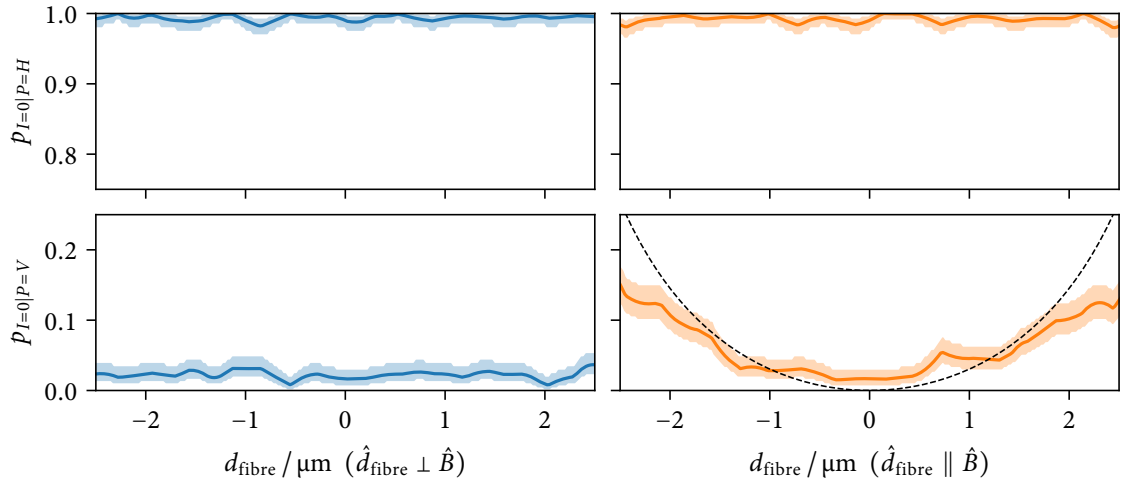
Figure 6.7: Ion state following detection of a photon of H and V polarisation (set up such that $\sigma^\pm \to H$, $\pi \to V$ on axis), given as the probability for the ion to fluoresce brightly at the end of the sequence (corresponding to $|0\rangle$). Shown are two slices through the data from fig. 6.6 around the point where the collection probability for π photons is maximised, perpendicular to the magnetic field direction (left), and along it (right). As the data was acquired on a grid in the laboratory coordinate system which is rotated approximately 45° compared to the magnetic field direction of interest here, the curves are bilinear interpolates between the data points (the grid points do not coincide exactly with the slices shown). The shaded bands give an approximate indication of the 1σ confidence interval expected from binomial statistics. The dashed line in the bottom-right slice depicts the theory prediction (§§4.2.3 and 4.3.1) for $\sigma^+ \to V$ polarisation mixing assuming nominal imaging system magnification, fibre mode field diameter, and piezo positioning stage calibration.
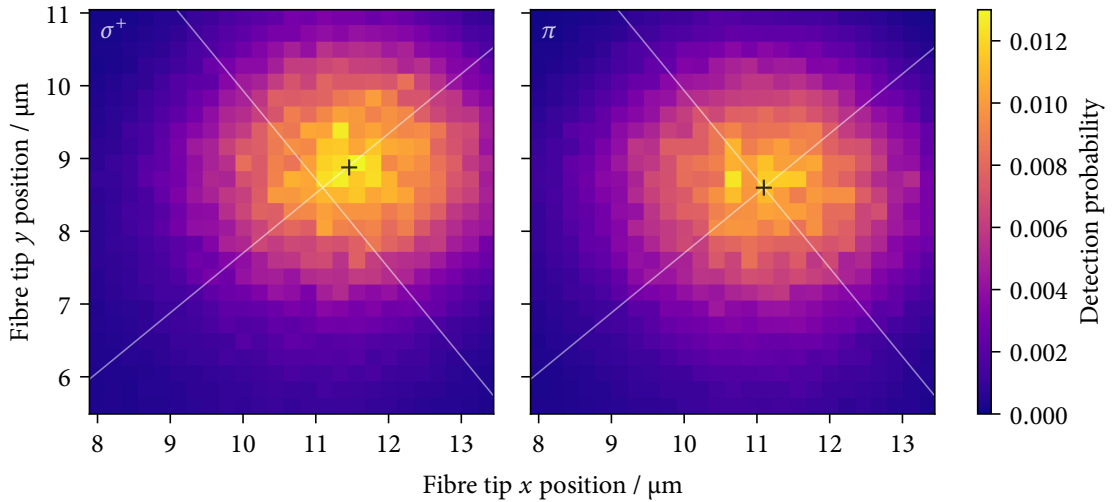
Figure 6.8: Overall collection efficiencies for photons from σ^+ (left) and π (right) decays. The x and y axes match the laboratory orientation shown in fig. 3.6 with the ambient magnetic field pointing diagonally in the down-right direction. The black plus signs mark the peak position inferred from a two-dimensional Gaussian fit to the data. The white crosses mark the directions parallel and perpendicular to the magnetic field (inferred from the $\sigma \to V$ polarisation mixing suppression data in fig. 6.6), centred on the π collection efficiency peak.

transition as a function of fibre position; an example for such data is shown in fig. 6.8. Compared to the steady-state fluorescence data in fig. 6.5, this cleanly shows only σ^+, not an equal mixture of $\sigma^\pm$, and also avoids the influence of $\sigma^\pm \to V$ polarisation mixing off-axis. The σ^+ collection efficiency peak is indeed shifted perpendicular to the magnetic field compared to the π peak, with the sign expected from the theory. The magnitude of the shift is 0.458(10) μm, which matches the theoretical prediction of 0.461 μm, though this close agreement might well be coincidence given that the analysis assumes the imaging system parameters and the voltage–distance transfer function of the piezoelectric translation stage to match the nominal values provided by the manufacturers.

In future work, it would be interesting to quantitatively probe the correspondence between the measurements presented here and the model from chapter 4. This will require some care, as the measurement is affected by drifts and hysteresis in the piezoelectric actuators used to position the fibre, other mechanical drifts in the system, the actual system magnification and travel of the piezoelectric translation stages per applied voltage (nominal values were used throughout), and any aberrations in the imaging system, such as caused by imperfect alignment.

6.2   Ion–photon entanglement

To study not just the purity of correlations between atomic decay channels and photon polarisation, but the full quantum state including coherences, the obvious tool is quantum state tomography – once the density operator describing the full state has been reconstructed, any further quantities can be derived from it.

Here, we will restrict the analysis to the state considered to the $\mathbb{C}^2 \otimes \mathbb{C}^2$ Hilbert space given by the Zeeman ground states on the ion and the polarisation degree of freedom on the photon, post-selected on the detection of a photon in one of the APDs. We perform all operations on the ion qubit in a rotating frame referenced to the laboratory frame at the time the photon is detected, after which we can ignore the continuous temporal degree of freedom, and the ion–photon state in the above discrete space should ideally be a maximally entangled two-qubit state (see §4.1).

The most straightforward method to reconstruct the state is to consider the ion state post-selected on the detection of a photon in a single detector, and vary both the pulses setting the ion readout basis and the wave-plate angles selecting the detected polarisation such that they span the complete Hilbert space (e.g. by measuring each qubit system along the three Cartesian coordinate axes of the Bloch sphere). The data acquisition for this proceeds as described in §6.1.1. The tomographic reconstruction requires some care, as a photon will only be detected with some total efficiency $0 < \eta < 1$. With only a single detector, we necessarily need to make the assumption that η is stable across the duration of the experiment (which is a very good approximation for the losses in the optical system

and detector efficiencies, but not necessarily for the pulsed laser excitation or the alignment of the collection fibre). Under this assumption, the ion–photon state can be reconstructed using a maximum likelihood estimation (MLE) technique, which I have first described in ref. [Nad16].

Concretely, let Γ_i be the projector corresponding to the i-th photonic measurement settings. For post-selection on APD0, which by convention corresponds to H polarisation (see fig. 3.7), this is

$$\Gamma_i = U_{\mathrm{wp}}(\alpha_4, \beta_{4,i})^\dagger\, U_{\mathrm{wp}}(\alpha_2, \beta_{2,i})^\dagger\, |H\rangle\langle H|\, U_{\mathrm{wp}}(\alpha_2, \beta_{2,i})\, U_{\mathrm{wp}}(\alpha_4, \beta_{4,i}), \tag{6.3}$$

where α_2 and α_4 are the actual retardances of the half- and quarter-wave plates, and $(\beta_{2,i}, \beta_{4,i})$ their respective orientations for the i-th photonic measurement settings. Let n_i the total number of photons observed for this detector setting (corresponding to the POVM element $\eta\Gamma_i$), and $N_i > n_i$ the total number of excitation attempt cycles necessary to do so[3]. Similarly, let Ξ_j be the projector corresponding to a bright ion measurement outcome for the j-th ion measurement setting. For measurements in the computational basis, $\Xi_j = |0\rangle\langle 0|$, or with an additional rotation,

$$\Xi_j = R_{xy}(\vartheta_j, \phi_j)^\dagger\, |0\rangle\langle 0|\, R_{xy}(\vartheta_j, \phi_j). \tag{6.4}$$

Let b_{ij}, d_{ij} be the number of bright and dark outcomes observed for the photon and ion measurement settings with the given indices ($\sum_j b_{ij} + d_{ij} = n_i$). After dropping some combinatorial pre-factors independent of the model parameters, the overall likelihood for a given set of observed events given an ion–photon density matrix ρ_{IP} and detection efficiency η is $\mathcal{L}(\rho_{IP}, \eta; (b_{ij})_{ij}, (d_{ij})_{ij}, (n_i)_i, (N_i)_i)$

$$= \prod_i \mathrm{tr}\left(\eta\Gamma_i \rho_P\right)^{n_i}\, \mathrm{tr}\left((\mathbb{1} - \eta\Gamma_i)\rho_P\right)^{N_i - n_i}$$

$$\cdot \prod_{ij} \frac{1}{\mathrm{tr}\left(\Gamma_i \rho_P\right)^{b_{ij} + d_{ij}}}\, \mathrm{tr}\left((\Xi_j \otimes \Gamma_i)\rho_{IP}\right)^{b_{ij}}\, \mathrm{tr}\left(((\mathbb{1} - \Xi_j) \otimes \Gamma_i)\rho_{IP}\right)^{d_{ij}}, \tag{6.5}$$

where $\rho_P = \mathrm{tr}_I \rho_{IP}$ is the reduced state of the photon. The maximum-likelihood estimate for the state is then

$$\tilde\rho_{IP}, \tilde\eta = \operatorname*{argmax}_{\rho_{IP}, \eta} \mathcal{L}(\rho_{IP}, \eta; (b_{ij})_{ij}, (d_{ij})_{ij}, (n_i)_i, (N_i)_i). \tag{6.6}$$

Compared to the original description in ref. [Nad16], we have not included any readout normalisation, as the Zeeman qubit state can efficiently be read out using the 674 nm laser

---

[3]We do not assume knowledge of the photon state in the analysis. In the experiment, we just choose a number of successful detections for each measurement setting, $n_i = \mathrm{const.}$, and record the number of attempts N_i until that number of successes is reached. Fixing the number of herald events slightly simplifies some aspects of our real-time data acquisition implementation and as the reduced state of the photons is approximately maximally mixed, $N_i \approx \frac{2n_i}{\gamma}$, we still allocate approximately the same amount of data acquisition time to each measurement basis. If the polarisation state was such that it contained very little of the post-selected state for some wave-plate settings, a strategy fixing the number of attempts instead would give more balanced statistical uncertainties.

Alice (left node):

	$\langle 0H|$	$\langle 1H|$	$\langle 0V|$	$\langle 1V|$
$|0H\rangle$	0.080			
$|1H\rangle$	$0.175 \cdot e^{i0.576\pi}$	0.431		
$|0V\rangle$	$0.173 \cdot e^{i0.332\pi}$	$0.403 \cdot e^{i1.751\pi}$	0.395	
$|1V\rangle$	$0.080 \cdot e^{i1.875\pi}$	$0.198 \cdot e^{i1.307\pi}$	$0.181 \cdot e^{i1.548\pi}$	0.094

Bob (right node):

	$\langle 0H|$	$\langle 1H|$	$\langle 0V|$	$\langle 1V|$	
0.499					$|0H\rangle$
$0.087 \cdot e^{i1.524\pi}$	0.023				$|1H\rangle$
$0.075 \cdot e^{i1.406\pi}$	$0.011 \cdot e^{i1.882\pi}$	0.012			$|0V\rangle$
$0.466 \cdot e^{i1.946\pi}$	$0.094 \cdot e^{i0.416\pi}$	$0.065 \cdot e^{i0.534\pi}$	0.466		$|1V\rangle$

Scale: 0.5, 0.25, 0.1, 0.01. Phase colour wheel: i, $+$, $-$, $-i$.
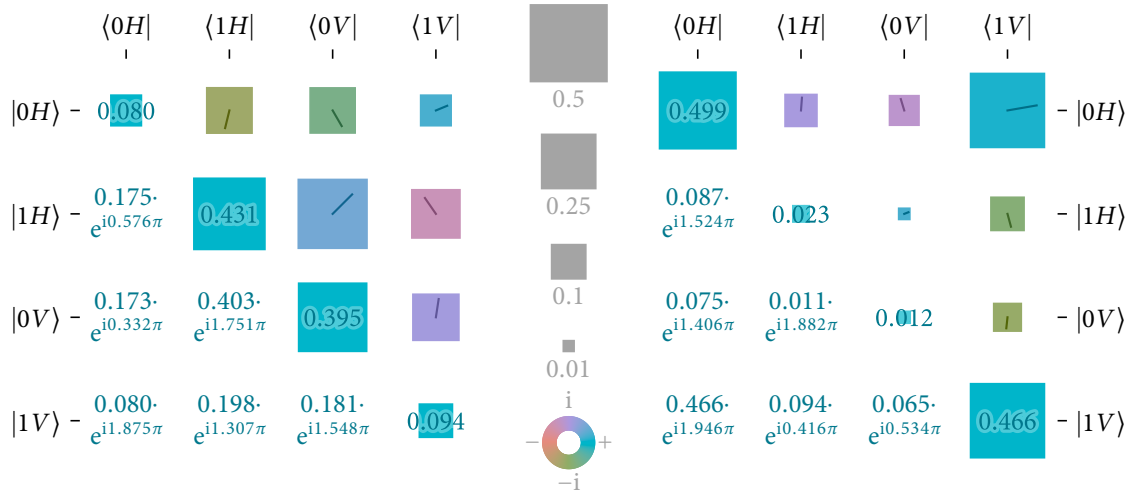
Figure 6.9: Maximum-likelihood estimates of the joint ion–photon state produced in nodes Alice (left) and Bob (right), obtained from 624 000 and 572 000 individual state tomography measurements, respectively. The states are rotated away from the ideal Bell state $|\Phi^+\rangle = (|0H\rangle + |1V\rangle)/\sqrt{2}$ due to the a priori unknown unitary action of the birefringence in the fibre carrying the photons.

pulses. The control system hardware now also supports deterministic real-time phase updates relative to the photon detection time, so instead of relying on the spread in photon arrival times to result in different phases shot-to-shot[4], we now just select between one of several settings as denoted by the index j.

For the experiment, we typically choose an intentionally over-complete set of measurement bases[5], given by the Cartesian product of $\beta_4 \in \left\{0, \frac{\pi}{4}\right\}$, $\beta_2 \in \left\{0, \frac{\pi}{8}, \frac{\pi}{4}, \frac{3\pi}{8}\right\}$,[6] and measuring the ion approximately $\frac{1}{3}$ of the time in the computational basis, the other $\frac{2}{3}$ rotated by $R_{xy}(\frac{\pi}{2}, \phi_k)$, $(\phi_k)_{k=0,\ldots,7} = \left(\frac{k\pi}{4}\right)_{k=0,\ldots,7}$, i.e. measuring along one of eight equally spaced axes in the xy plane of the ion Bloch sphere. The order of points can be randomised to avoid slow collection efficiency drifts resulting in systematic biases (although we typically find that this is not necessary given the high data acquisition rates, preferring to rotate the motorised wave plates less to avoid the associated execution time overhead). We also typically apply a spin echo (Hahn echo) to the ion qubit, where initially the ion state is just inverted using the 674 nm laser and mapped back to the Zeeman qubit before the $\frac{\pi}{2}$ analysis pulse, though this became less important following improvements to the technical delay between herald event and analysis sequence, and the ambient magnetic field stability.

The results for two such tomography runs, one in each of the nodes Alice and Bob, are

---

[4]This approach was a natural fit for the experiments in ref. [Nad16], where the higher magnetic field of 14.6 mT led to a Zeeman splitting much larger than the excited state linewidth.

[5]Such an over-complete set of measurements is useful to verify the calibration of the wave plate retardances (which can be difficult to establish accurately if not done in place), or for the investigation of sign issues between theory and implementation.

[6]For ideal wave plates, several of these angle combinations would give rise to identical measurement projectors.

Alice:

	$\langle 0H|$	$\langle 1H|$	$\langle 0V|$	$\langle 1V|$
$|0H\rangle$	0.472			—
$|1H\rangle$	$0.010\cdot e^{i1.114\pi}$	0.003		
$|0V\rangle$	$0.009\cdot e^{i0.833\pi}$	$0.001\cdot e^{i1.740\pi}$	0.003	
$|1V\rangle$	$0.482\cdot e^{i2.000\pi}$	$0.008\cdot e^{i0.636\pi}$	$0.008\cdot e^{i1.299\pi}$	0.522

Bob:

	$\langle 0H|$	$\langle 1H|$	$\langle 0V|$	$\langle 1V|$	
0.510			—	$|0H\rangle$	
$0.008\cdot e^{i0.971\pi}$	0.004			$|1H\rangle$	
$0.010\cdot e^{i1.008\pi}$	$0.002\cdot e^{i1.950\pi}$	0.001		$|0V\rangle$	
$0.478\cdot e^{i2.000\pi}$	$0.016\cdot e^{i0.955\pi}$	$0.014\cdot e^{i1.041\pi}$	0.485	$|1V\rangle$	

Scale legend: 0.5, 0.25, 0.1, 0.01; phase: i, $+$, $-$, $-i$.
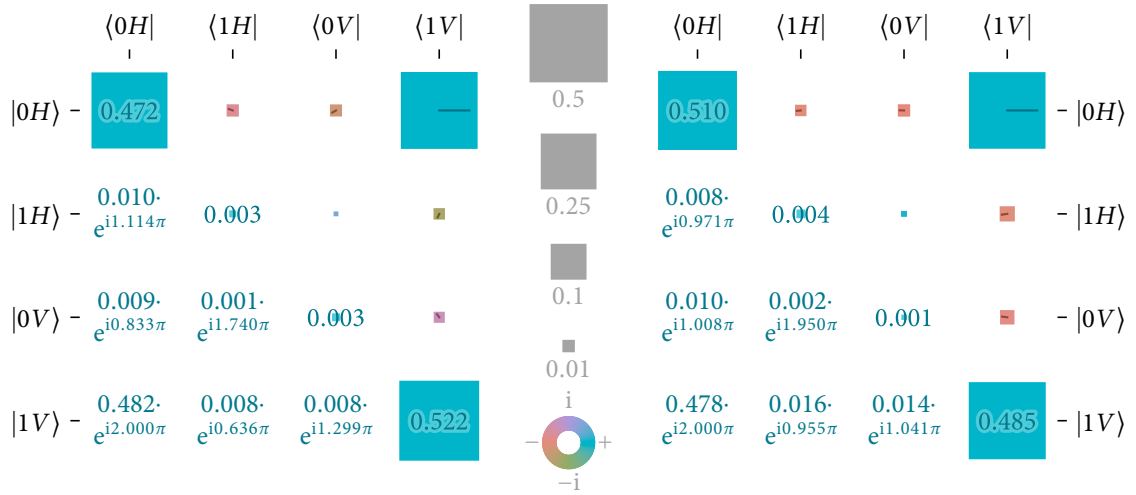
Figure 6.10: Maximum-likelihood ion–photon state estimates in nodes Alice (left) and Bob (right) from fig. 6.9, transformed post-hoc with unitary operations acting locally on the two subsystems to maximize the overlap with the Bell state $|\Phi^+\rangle = (|0H\rangle + |1V\rangle)/\sqrt{2}$, which corrects for the fibre birefringence and makes it visually clear that the states are in fact close to a maximally entangled state.

System	Alice	Bob
Fully entangled fraction $\mathcal{F}$	0.979(3)	0.977(1)
Entanglement of formation E_F	0.941(8)	0.938(3)
Logarithmic negativity $\mathcal{N}$	0.969(4)	0.966(1)
Single-detector efficiency	0.403 %	0.568 %

Table 6.1: Ion–photon state figures of merit, derived from the MLE tomography estimate (errors are the standard deviation obtained from bootstrapping). The efficiency is given end-to-end from entanglement trials to clicks on APD0; the combined efficiency for all four detectors is approximately a factor of 4 higher. The collection lens alignment in node Alice had degraded at the time this dataset was obtained, giving rise to the lower efficiency.

shown in fig. 6.9. For these data, 624 000 copies of the ion–photon state were measured in Alice, 572 000 in Bob, distributed over 27 min and 18 min of wall-clock time, respectively. The density matrix estimates were obtained by a numerical implementation of eq. (6.6) in the form of a direct minimisation of the negative log-likelihood $-\log(\mathcal{L}(\rho_{IP}, \eta; \dots))$ according to eq. (6.5). This optimisation has $15+1$ real degrees of freedom parametrising the state and the overall collection efficiency. To ensure positivity of the density matrix, the matrix was expressed in its Cholesky decomposition $\rho = LL^\dagger$, where L is a lower-triangular matrix with non-negative diagonal entries, the entries of which are optimised over. A summary of several figures of merit is shown in table 6.1, along with an estimate for the standard errors obtained using bootstrapping.

The form of the resulting states depends on the uncontrolled unitary transformation effected by the polarisation rotation in the optical fibres, as well as the choice of phase reference for the ion Zeeman qubit manipulation, which depends on the absolute latency
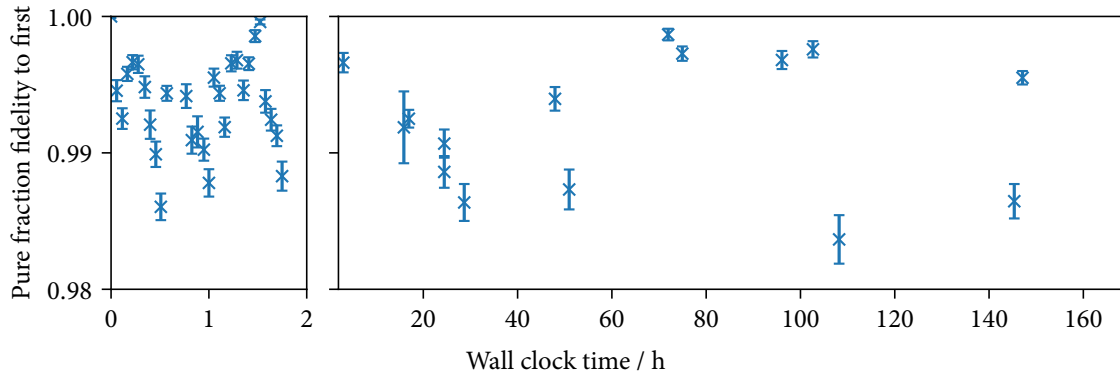
Figure 6.11: Drift in the ion–photon entangled state reconstructed from tomography measurements across one week (node Alice, for the states post-selected on a click on APD0, with spin echo on the ion qubit). To decouple the overlap data from entanglement imperfections ($\mathcal{F} \sim 96\,\%$ for the data presented here), the fidelity of the largest pure state component (density matrix eigenvector) at each point to that of the first measured state is shown, in form of the 1σ confidence interval obtained from bootstrapping. The drift is caused by birefringence variations in the optical fibre links, whose magnitude depends on the thermal stability and any mechanical disturbances; oscillations on the time scale typical for temperature changes due to the laboratory air conditioning are visible in the first, more densely sampled section.

through the photon detection and RF synthesis chain. For ease of visual interpretation, the same MLE results are also shown in fig. 6.10 post-processed using qubit-local rotations to maximise their overlap with $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|0H\rangle + |1V\rangle)$. In practice, we also use a similar optimisation procedure on the tomography data to initialise the wave plate settings for the ion–ion entanglement use case, which we choose such that $p_{I=0|P=H}$ is maximal. Drifts in birefringence due to ambient temperature and mechanical stresses lead to slow changes in the unitary polarisation rotation induced by the fibre depending on environmental conditions, but we typically observe the state to be stable to $\geq 98\,\%$ over days (see fig. 6.11).

To improve data acquisition rates while continuing to use the ion–ion entanglement generation setup (Bell-state analyser) as the polarisation-sensitive detector, it would be desirable to extend this tomography technique to make use of photons in all four APD detectors. This is relatively straightforward but requires the introduction of some knowledge about the exact behaviour of the Bell-state measurement apparatus into the device model: as the attempt loop will terminate after the detection of a photon in *any* of the APDs, the data cannot simply be divided up into four streams of the form described here after the fact. A description of such as a scheme which is now routinely used in the experiment will be given in P. Drmota's upcoming DPhil thesis; the tomography data presented here is older and was still acquired using only one detector.

6.3   Ion–ion entanglement

The simultaneous detection of two photons of different polarisation in the Bell-state analyser heralds the creation of an approximately maximally-entangled ion–ion state. As dis-

cussed in §4.4, there are four different heralding detector click patterns, which in the ideal case leave the ions in a pure state independent of the photon detection times, given by

$$\left|\Psi_{AB}^{\pm}\right\rangle = \frac{1}{\sqrt{2}}\left(\left|0_A 1_B\right\rangle \pm e^{i\phi}\left|1_A 0_B\right\rangle\right),\tag{6.7}$$

where ϕ is an a priori unknown phase set by the residual fibre birefringence (along the third axis not compensated by the wave plates), and in practice also the choice of phase reference for the manipulation of the Zeeman qubits. The states for the cases where the photons are detected on the same side of the 50 : 50 beamsplitter are predicted to be the same, and so are those where detectors on opposite sides click, corresponding to the two signs in eq. (6.7). This does not hold in the presence of various technical imperfections, so in the following, we treat each heralding detector click pattern separately.

To analyse the creation of entanglement in the experiment, we again follow the blueprint from §6.1.1, now synchronised between the two ion trap nodes: we (re)initialise the ions in $\left|0\right\rangle$, excite them, and observe APD detector clicks until we detect an acceptable herald pattern. The total success probability now goes with the square of the single ion–photon efficiency, so is on the order of 10^{-4}. For the tomographic reconstruction of the created states, we either measure directly in the computational basis after mapping to the $\left|0\right\rangle \coloneqq \left|5s\,^2S_{1/2}, m_J = -\frac{1}{2}\right\rangle \leftrightarrow \left|1'\right\rangle \coloneqq \left|4d\,^2D_{5/2}, m_J = -\frac{3}{2}\right\rangle$ optical qubit, or apply extra $\pi/2$ pulses to measure along the other two axes of the respective Bloch sphere. In other words, for each of the 3×3 combinations, we collect statistics for a projective measurement onto one of four Pauli matrix tensor product eigenstates.

As the ideal state does not evolve in time (except for a global phase), the choice of phase origin for these measurements does not matter, as long it is synchronised between the ions; we choose the end of the last (successful) 1 μs attempt cycle as the reference time for the sake of it being well-defined across the systems. No spin echo is performed, as it does not further improve the state fidelity for the short ($\sim 30\,\mu s$) wait times between successful herald and mapping into the optical qubit. To suppress the effect of non-ideal "map" π pulses, such as due to motional heating across the attempt sequence, an additional shelving π pulse $\left|5s\,^2S_{1/2}, m_J = \frac{1}{2}\right\rangle \to \left|4d\,^2D_{5/2}, m_J = \frac{3}{2}\right\rangle$ is applied instead of the $\pi/2$ pulse for measurements in the computational basis (though not for measurements in the X and Y bases, as this cannot preserve the coherence between the states). We sequentially measure in each combination of bases for a predetermined number of successful heralds, accepting all four detector click patterns every time and dividing the data into separate streams for each pattern after the fact. As the four patterns occur with roughly equal probabilities – up to differences in the detector efficiencies –, the statistical uncertainties in the estimates will be roughly matched between the patterns nevertheless.

A first set of measurements, taken in April 2019, was reported in ref. [Ste+20], with an average fidelity of $\mathcal{F} = 94.0(5)\,\%$ and rate of $182\,\text{s}^{-1}$. It was analysed the same way as discussed here for the dataset from table 6.2, which was acquired more recently, in May

O_A	O_B	$+_A+_B$	$+_A-_B$	$-_A+_B$	$-_A-_B$
X	X	229	332	375	284
X	Y	27	526	650	15
X	-X	317	184	310	374
X	-Y	532	23	24	580
X	Z	541	558	650	609
Y	X	548	16	20	579
Y	Y	284	301	389	241
Y	-X	25	509	618	27
Y	-Y	286	253	260	379
Y	Z	551	626	715	594
-X	X	356	296	237	347
-X	Y	618	19	27	500
-X	-X	271	362	321	211
-X	-Y	20	633	563	29
-X	Z	684	671	539	514
-Y	X	23	621	527	17
-Y	Y	400	231	286	292
-Y	-X	617	33	23	516
-Y	-Y	251	378	308	277
-Y	Z	717	594	544	561
Z	X	525	693	620	527
Z	Y	655	562	634	502
Z	-X	661	579	635	589
Z	-Y	543	674	521	575
Z	Z	82	2415	2378	17

(a) Number of observations per Pauli eigenstate for detector pattern 0011, for a total of 43 182 clicks.

O_A	O_B	$+_A+_B$	$+_A-_B$	$-_A+_B$	$-_A-_B$
X	X	229	358	408	305
X	Y	36	553	667	16
X	-X	363	229	322	417
X	-Y	558	35	34	698
X	Z	581	631	774	669
Y	X	589	28	32	646
Y	Y	286	343	470	259
Y	-X	20	593	679	26
Y	-Y	346	284	255	447
Y	Z	560	681	731	620
-X	X	409	327	235	354
-X	Y	759	26	29	551
-X	-X	283	433	370	236
-X	-Y	23	662	594	31
-X	Z	773	755	580	643
-Y	X	39	645	582	22
-Y	Y	411	271	288	351
-Y	-X	691	33	28	609
-Y	-Y	277	461	301	282
-Y	Z	707	674	556	587
Z	X	640	713	620	605
Z	Y	735	568	668	605
Z	-X	679	653	609	604
Z	-Y	550	828	671	662
Z	Z	69	2603	2517	13

(b) Number of observations per Pauli eigenstate for detector pattern 1100, for a total of 47 278 clicks.

O_A	O_B	$+_A+_B$	$+_A-_B$	$-_A+_B$	$-_A-_B$
X	X	357	242	310	368
X	Y	588	21	14	658
X	-X	247	338	346	357
X	-Y	19	571	658	29
X	Z	596	608	683	659
Y	X	19	603	680	26
Y	Y	284	267	277	375
Y	-X	604	18	28	643
Y	-Y	289	284	408	273
Y	Z	506	716	695	658
-X	X	345	359	329	251
-X	Y	20	674	565	27
-X	-X	352	340	292	325
-X	-Y	637	17	26	573
-X	Z	694	668	562	562
-Y	X	685	20	28	560
-Y	Y	237	383	335	300
-Y	-X	28	623	604	15
-Y	-Y	372	298	287	343
-Y	Z	723	629	534	735
Z	X	797	573	556	659
Z	Y	608	727	629	642
Z	-X	548	785	697	524
Z	-Y	667	623	657	569
Z	Z	50	2489	2513	24

(c) Number of observations per Pauli eigenstate for detector pattern 1010, for a total of 46 016 clicks.

O_A	O_B	$+_A+_B$	$+_A-_B$	$-_A+_B$	$-_A-_B$
X	X	337	201	296	369
X	Y	571	40	22	596
X	-X	195	373	343	285
X	-Y	35	548	633	23
X	Z	564	575	664	638
Y	X	31	562	589	32
Y	Y	341	257	208	418
Y	-X	556	29	50	575
Y	-Y	230	356	396	254
Y	Z	464	662	681	540
-X	X	242	356	365	192
-X	Y	23	608	525	29
-X	-X	365	248	204	387
-X	-Y	618	19	23	532
-X	Z	622	622	556	555
-Y	X	618	30	22	561
-Y	Y	246	410	330	229
-Y	-X	41	579	541	19
-Y	-Y	370	236	224	335
-Y	Z	762	537	560	580
Z	X	708	515	583	666
Z	Y	526	704	588	647
Z	-X	481	736	691	529
Z	-Y	706	545	597	612
Z	Z	60	2284	2459	27

(d) Number of observations per Pauli eigenstate for detector pattern 0101, for a total of 43 524 clicks.

Table 6.2: Number of observations of each eigenstate for tensor product combinations $O_A \otimes O_B$ of Pauli operators in the Alice-Bob state tomography experiment, for each of the four detector click patterns.
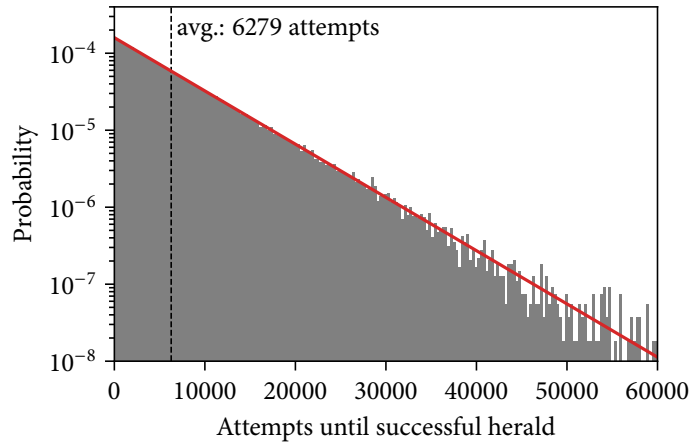
Figure 6.12: Number of shots before herald for the ion–ion state tomography dataset shown in table 6.2. With an effective attempt rate of 625 kHz (including periodic cooling), the entanglement generation rate is $100\,\mathrm{s}^{-1}$.

2021, following a number of experimental upgrades that greatly increased automation and robustness[7]. It consists of 180 000 total heralds spread across the four detector click patterns. To avoid any potential bias in the xy plane introduced by gate rotation angle errors (cf. §7.4.1), we also include, in addition to the measurement settings $\{X, Y, Z\}$, the opposite basis change rotations leading to $-X$ and $-Y$ (which, ideally, just lead to mirrored outcomes).

The distribution of excitation attempts before observing a herald event for this dataset is shown in fig. 6.12. A 30 ns coincidence acceptance window aligned with the excitation laser pulse was used to capture most events while excluding background counts ($\approx 98\,\%$, although this could likely be tuned further). As both the pulsed laser excitation efficiencies and fibre-coupling efficiencies had degraded due to misalignments in the optical paths compared to the April 2019 data in ref. [Ste+20], the success probability was only 1.588×10^{-4}. Additionally, an attempt timeout of only 500 μs had been used for this dataset by accident, leading to an overall rate of only $100\,\mathrm{s}^{-1}$. This is solely a consequence of the data having been acquired at the end of a long effort focussed on automation and long-term robustness, not peak performance; there is no reason to expect the $182\,\mathrm{s}^{-1}$ rate not to be restored (and exceeded) at the same higher level of fidelity following realignment of the collection optics.

The observation counts from table 6.2, along with the corresponding ideal projectors, make up the entire input for the state tomography procedure (i.e., no extra readout correction, etc. is employed). The likelihood function for the number of times n_k each measurement outcome k was observed, given the ion–ion density matrix ρ_{AB}, is simply (up to a

---

[7]This particular dataset was chosen as it was contemporaneous to the DIQKD experiments described in the next chapter, but is representative of the experiment performance at the time.

Detector click pattern	0011	1100	0101	1010
Fully entangled fraction $\mathcal{F}$	0.958(2)	0.958(2)	0.961(3)	0.964(2)
Entanglement of formation E_F	0.887(7)	0.885(7)	0.890(8)	0.899(6)
Logarithmic negativity $\mathcal{N}$	0.938(3)	0.938(4)	0.943(4)	0.948(3)
Herald probability	3.82×10^{-5}	4.18×10^{-5}	3.85×10^{-5}	4.07×10^{-5}

Table 6.3: Commonly cited figures of merit quantifying the amount of entanglement in the maximum-likelihood density matrix estimates.

combinatorial factor independent of the state)

$$\mathcal{L}(\rho_{AB}; (n_k)_k) = \prod_{k=1}^{24\times4} \mathrm{tr}(E_k \rho_{AB})^{n_k}, \tag{6.8}$$

where E_k are the projectors onto the various measurement outcomes. The MLE state estimate that maximises eq. (6.8) can be obtained by any number of optimisation methods; I chose the diluted fixed-point iteration algorithm described by Reháček et al. [Reh+07] (which, as any sensible MLE procedure, ensures that the result is physical, i.e. Hermitian, positive-semidefinite and of unit trace). The resulting density matrices are shown in fig. 6.13. Table 6.3 summarises a few figures of merit calculated from the MLE predictions. Most commonly reported in the literature on remote entanglement generation is the fidelity to the closest Bell state, or maximally entangled fraction; $\mathcal{F} \approx 96\,\%$ is observed across all four heralding detector click patterns here.

The average performance, $\mathcal{F} = 96.0(1)\,\%$ at a rate of $100\,\mathrm{s}^{-1}$, currently represents the state of the art across matter qubit remote entanglement experiments (see §6.4). Thus, the statistical confidence in the fidelity estimate merits particular attention.

6.3.1 Confidence intervals

Maximum-likelihood tomography is widely used in practice, as it is conceptually simple, has a clear statistical interpretation, and performs well on the nearly-pure states that are often of interest. Sampling noise, however, tends to be non-negligible in real-world applications – not least, one would not want to spend more laboratory time than necessary gathering the data –, such that a point estimate alone is not very helpful: What one would really like to obtain is some form of confidence region describing the range of values the true quantities are likely to inhabit.

One commonly used approach to this is *bootstrapping*, that is, generating many artificial datasets with the same statistical properties as the experimental results and analysing them the same way. The statistical uncertainty in some derived quantity of interest is then estimated by recording its distribution over the artificial tomography results. This is simple to implement numerically and works in many different situations without requiring
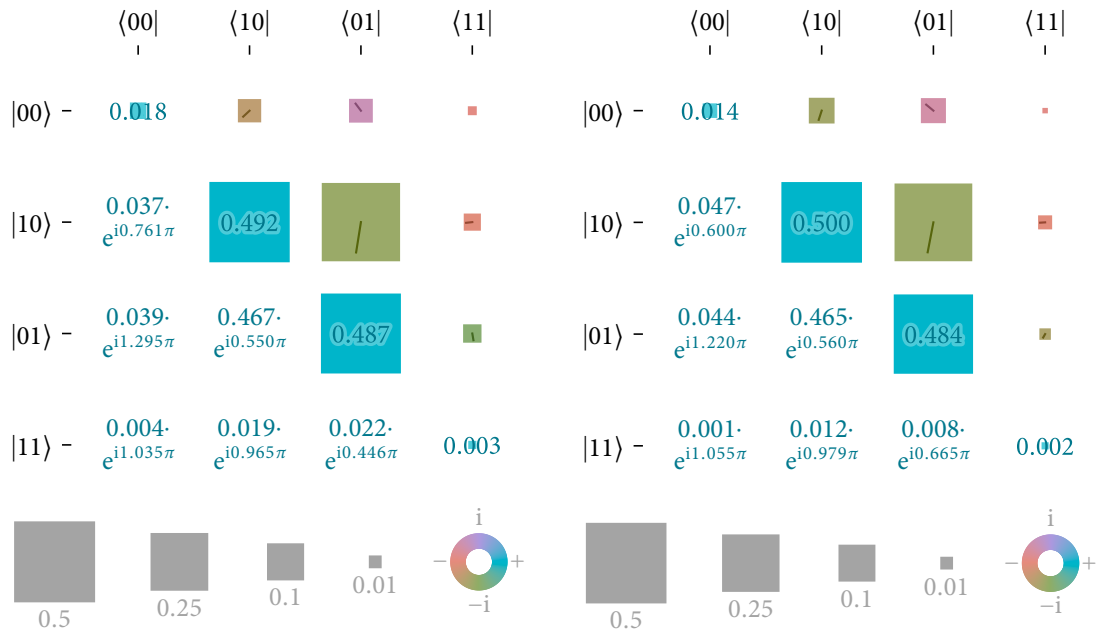
(a) ρ_{AB} for detector click pattern 0011, estimated from a total of 43 182 clicks.

(b) ρ_{AB} for detector click pattern 1100, estimated from a total of 47 278 clicks.
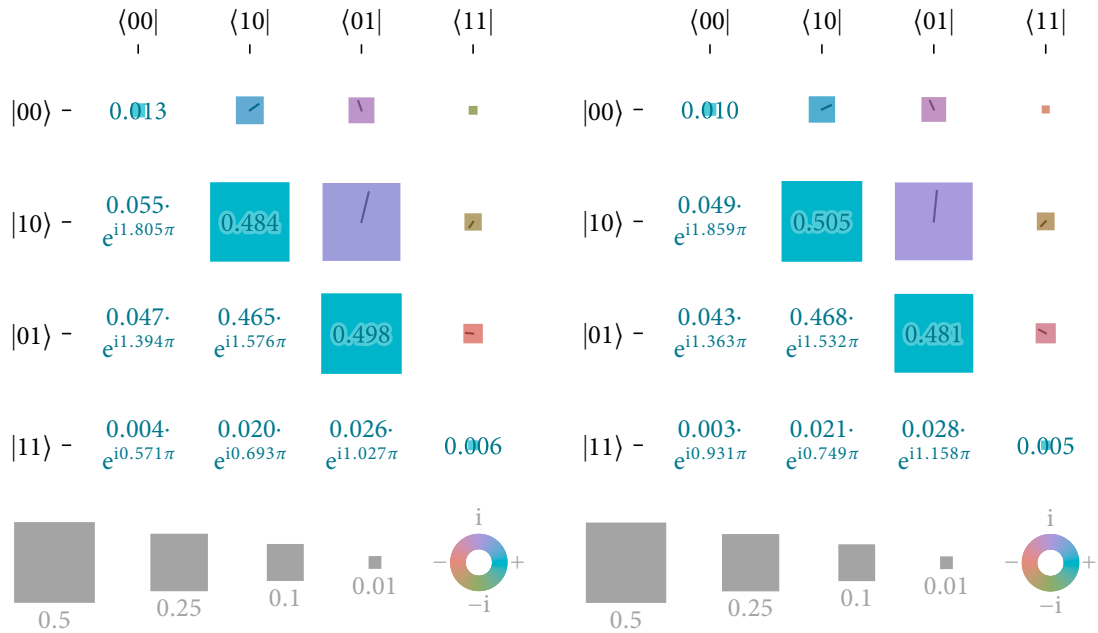
(c) ρ_{AB} for detector click pattern 0101, estimated from a total of 43 524 clicks.

(d) ρ_{AB} for detector click pattern 1010, estimated from a total of 46 016 clicks.

Figure 6.13: Density matrices for the joint states ρ_{AB} of Alice's and Bob's ions following each of the four heralding detector click patterns, as obtained by maximum-likelihood estimation for the data from table 6.2, representing 180 000 total herald events across all the detector click patterns. See table 6.1 for numerical figures of merit derived from this.

extra theoretical work, but precise statements about the meaning of its results are hard to make (at least if they are required to be useful).

Using the MLE density matrix estimate for the calculation of derived quantities is a perilous undertaking also for other reasons. For instance, MLE can be shown not to be optimal (in the sense of admissibility) for commonly used error metrics [FB18], and there is work discussing bias in derived quantities [SGV17] and implausibility of its results [Blu10]. In particular, the fact that the estimators are biased leads to bootstrapped error bars not necessarily being a good indicator for the uncertainty in the state estimate, as discussed in ref. [Blu12].

Alternative methods to obtain estimates from tomography data along with some measure for their uncertainty exist, such as various Bayesian approaches, but they are not without their own challenges and shortcomings. At this point, the field has not adopted a single solution as the "gold standard" yet. Thankfully, the two-qubit case is reasonably cheap to tackle in terms of computational effort, such that we can easily implement and compare various approaches directly, here taking the fully entangled fraction as the quantity of interest. Concretely, we will compare two bootstrapping methods with a more principled approach to constructing error bars from likelihood-weighted density matrix samples (which also has a Bayesian interpretation).

For bootstrapping, there are two main variants in the context of state tomography, often referred to as parametric and non-parametric bootstrapping. In the parametric case, the MLE density matrix estimate is used as the model, and further instances of the experiment simulated using a multinomial outcome distribution according to the ideal Pauli-product measurement projectors. The distribution of the quantities of interest across the tomographic reconstructions for each of these synthetic experiments is then used to derive e.g. an estimate for the standard error. For the non-parametric case, the synthetic experiments are generated by directly resampling the experimental outcomes, that is, by sampling from a multinomial distribution with probabilities given by the respective fractions of observed outcomes in the experimental dataset

Lastly, I implement a more principled approach in the spirit of the confidence regions proposed by Christandl and Renner [CR12]. The main idea here is that the likelihood function can be used to define a probability distribution μ by

$$\mu(\rho) := \frac{1}{c} \mathcal{L}(\rho; (n_k)_k), \tag{6.9}$$

where c is a normalisation factor depending on the choice of projectors such that $\int \mu(\rho) \, d\rho = 1$ for the Hilbert-Schmidt measure $d\rho$ (with $\int d\rho = 1$). The reduced distribution for a given derived figure of merit F is then

$$\mu_F(f) = \int \delta(F(\rho) - f) \, \mu(\rho) \, d\rho, \tag{6.10}$$

which we can approximate numerically following ref. [FR16] by sampling density matrices according to (6.9) using a Markov-chain Monte-Carlo algorithm (Metropolis–Hastings).
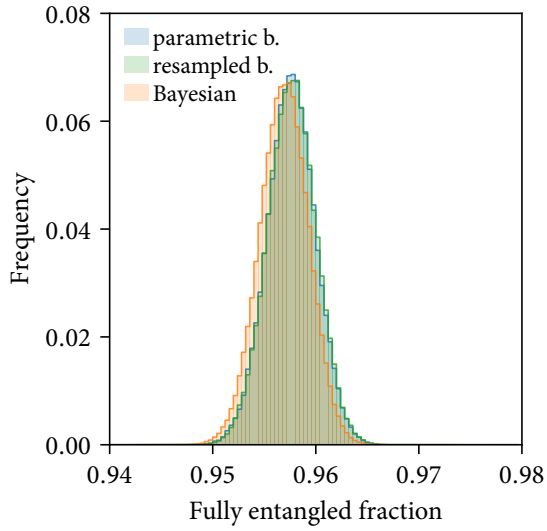
(a) Estimate density for detector click pattern 0011.
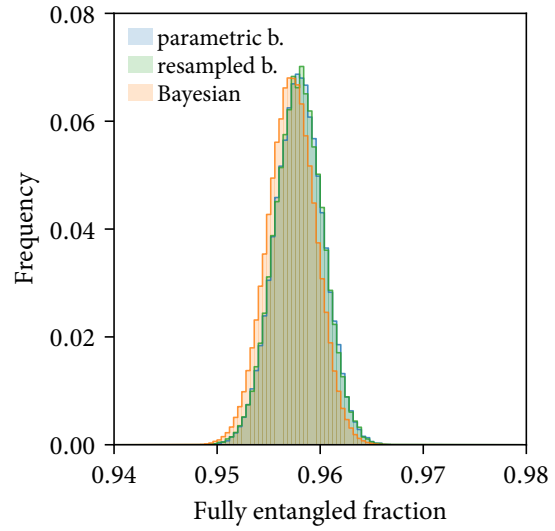
(b) Estimate density for detector click pattern 1100.
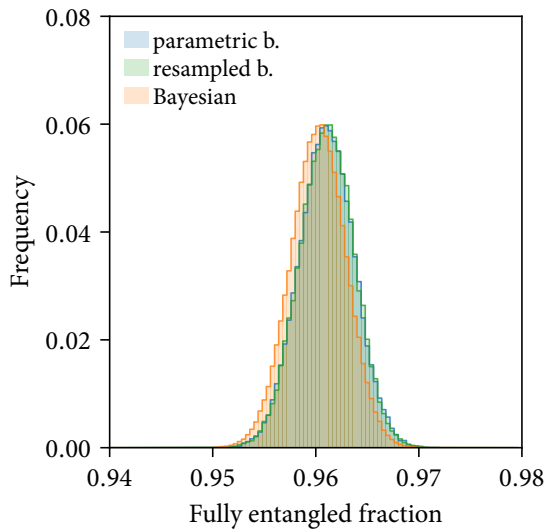
(c) Estimate density for detector click pattern 0101.
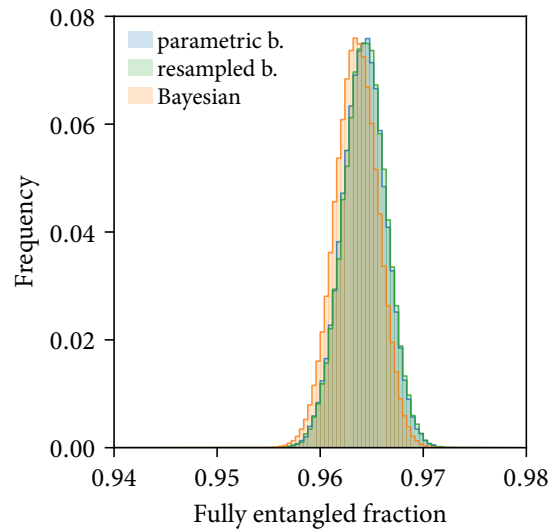
(d) Estimate density for detector click pattern 1010.

Figure 6.14: Distribution of estimates for the fully entangled fraction $\mathcal{F}$ (fidelity to closest maximally entangled state). Shown are histograms obtained by different methods for numerically estimating the figure-of-merit distributions: *parametric b.* refers to parametric bootstrapping using outcomes simulated from the MLE density matrix estimate, *resampled b.* to non-parametric bootstrapping by directly sampling from the observed outcomes with replacement, and *Bayesian* to integrating over likelihood-weighted samples obtained using the Metropolis–Hastings algorithm, which coincides with the posterior distribution for a Hilbert–Schmidt-uniform prior in a Bayesian picture.

Like for bootstrapping, the quantities $F(\rho)$ are collected into a histogram. In a Bayesian picture, μ coincides with the Bayesian posterior for a Hilbert–Schmidt-uniform prior, as proposed e.g. in ref. [Blu10]. This numerical approximation of $\mu_F(f)$ can also be used to construct confidence regions with precisely defined meaning, although I don't attempt this here as they tend to be unhelpfully large [CR12; FR16].

Figure 6.14 shows the results of applying all the above schemes to the fully entangled fraction (fidelity to the closest maximally entangled state). As before, the four herald patterns are analysed separately. The parametric and simple resampling bootstrap approaches give virtually identical results. As the data is decently sampled and there don't seem to be any gross irregularities, this isn't surprising. The fidelity distribution obtained using the Metropolis–Hastings samples is of approximately the same width, but is centred around slightly lower values. Such a bias was also observed in ref. [FR16], presumably as the volume factor increases towards less pure states. For the final values of $\mathcal{F}$ quoted, I chose to report the MLE point estimates together with the standard deviation of the estimate densities in accordance with the general practice in the field.

Density matrix rank. For all four of the maximum-likelihood estimates shown in fig. 6.13, the smallest singular value (i.e. magnitude of the smallest eigenvalue) is $< 2 \times 10^{-3}$. The tendency of MLE to produce reduced-rank estimates in the face of sampling noise or non-ideal data is well-documented [Blu10]; in many cases this can be understood as stemming from an actual likelihood maximum outside the set of physical (positive) density matrices, where the closest physical state is then rank-deficient. However, if the physical state is indeed close to being rank-deficient, this is of interest for entanglement distillation, as the corresponding type of error could then be amplified with little consequence. In our remote entanglement case this might be plausible, as neither $\sigma^+ \to V$ polarisation mixing, dark-to-bright readout errors (from imperfect mapping pulses or decay from the shelf level), nor ion dephasing lead to any population in $|11\rangle\langle11|$.

Figure 6.15 shows the results from sampling density matrices weighted by likelihood to estimate the distribution from eq. (6.10) for their three smallest singular values (using the same Metropolis–Hastings approach employed for the fidelity error estimates in fig. 6.14). This likelihood-weighted (or uniform-Bayesian-prior) analysis lends some credibility to the idea that a specialised distillation protocol that trades off a reduced input Bell pair count against an amplification of the errors in one particular channel might indeed work well to further increase the entangled state fidelity.
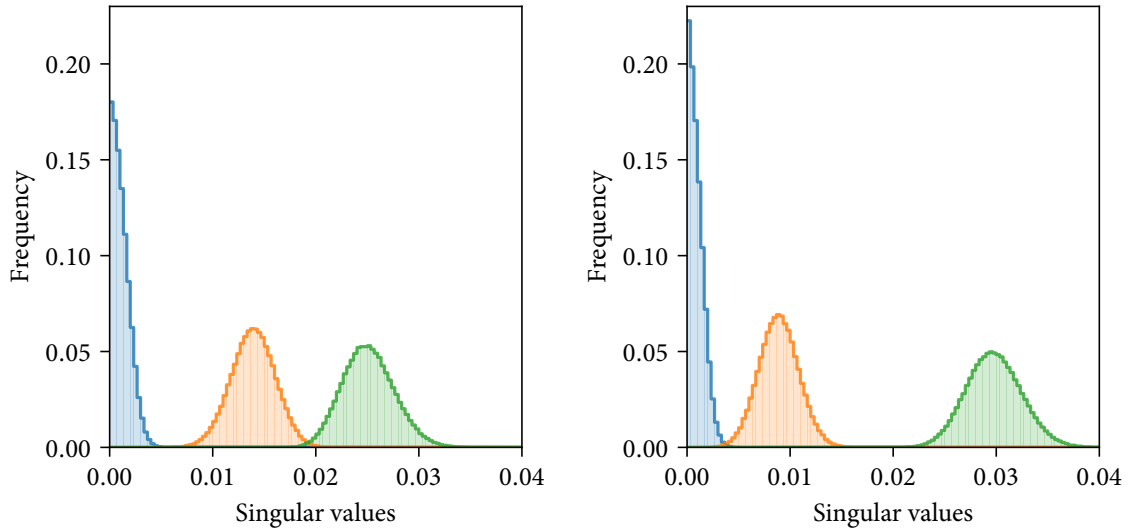
(a) Estimate density for detector click pattern 0011.

(b) Estimate density for detector click pattern 1100.
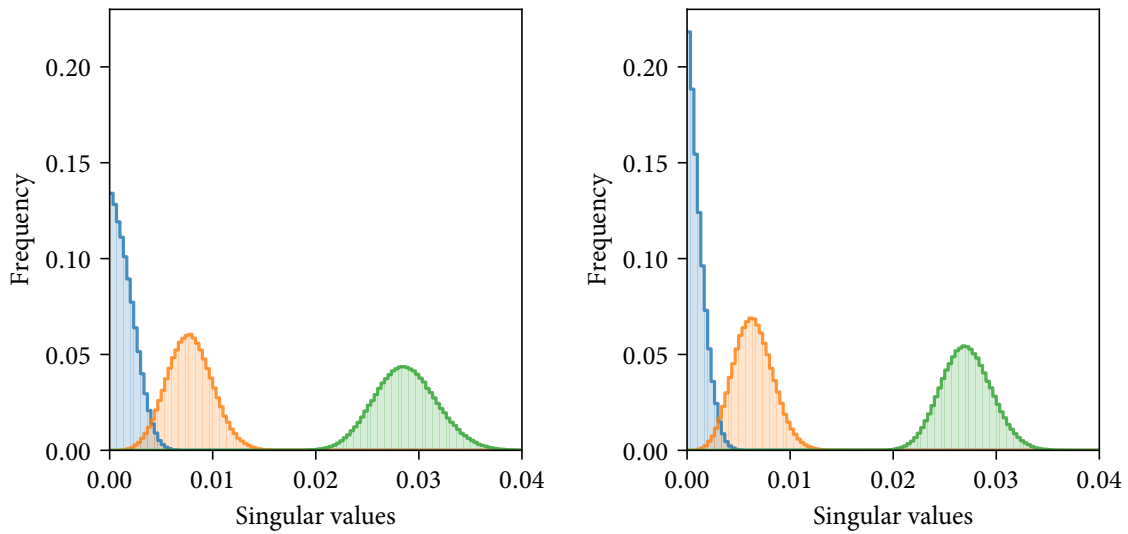
(c) Estimate density for detector click pattern 0101.

(d) Estimate density for detector click pattern 1010.

Figure 6.15: Distribution of estimates for the three smallest singular values of the ion–ion state ρ_{AB}. Shown are histograms obtained from sampling density matrices distributed according to their likelihood given the experimental tomography. Each bin is $\frac{1}{3} \cdot 10^{-3}$ in width; their sum is normalised to 1.

6.4 Discussion

6.4.1 Error sources

As our recent focus was on enabling the use of the remote entanglement link as a building block in more complex experiments (such as that described in chapter 7), with consequent demands on stability and automation, the construction and validation of a model for the remaining imperfections has not been a high priority – at least not before severe degradation of the vacuum in one of the nodes made further experiments impossible. The discussion here must thus remain rather perfunctory in nature.

The observations in §6.1 qualitatively validate the considerations from chapter 4; the correlations are considerably better than would be expected for collection with NA 0.6 in free space, and both the dependence of collection efficiencies and $\sigma^+ \rightarrow H$ polarisation mixing on the fibre position are as expected from theory. However, even with the fibre well-aligned, an imbalance in the errors $p_{I=0|P=V}$ and $p_{I=1|P=H}$ remains. As configured here, both polarisation mixing and errors in the 674 nm single-qubit gates (e.g. due to the thermal state of motion of the ions) would only produce $p_{I=1|P=H}$, so an imbalance of this sign is not implausible[8]. A likely explanation for residual polarisation mixing errors would be non-uniform birefringence in the high-NA portion of the optical system, for example from stresses on the vacuum viewport.

This imbalance in errors in the $\{|0\rangle, |1\rangle\}$ and $\{|H\rangle, |V\rangle\}$ bases is also visible in the ion–photon full tomography density matrices from §6.2. The reconstructed states furthermore suggest that there might be a slight additional loss in fidelity from an unidentified dephasing process (beyond that explained by ion coherence times and photon detection jitter, both of which are small). All in all, the $\approx 2\%$ (each) of ion–photon entanglement errors constitute the largest share of the 4% ion–ion infidelity[9]. The known imperfections of the Bell-state analyser (see fig. 3.7) only contribute $\approx 3 \times 10^{-3}$ to the ion–ion state error, although it would be interesting to cross-validate this, e.g. by monitoring the number of Hong–Ou–Mandel-forbidden coincidences. The framework described in chapter 4 is comprehensive and would readily support the matching of the experimental observed density matrices to models for any of these imperfections; this remains for future work, however.

Figure 6.16: Survey of literature on entanglement generation between remote matter qubits over an optical link, for which the entangled state generation rate and fidelity were either reported or possible to infer. Colours mark the kind of physical platform used: single ions trapped in RF traps, nitrogen-vacancy centres (colour centres) in diamond, single neutral atoms, and semiconductor quantum dots. Only qubit systems that allow manipulation of the stored information are shown, not pure quantum memories (with the exception of one point, where a single trapped atom was entangled with a collective excitation stored in a Bose–Einstein condensate). Symbols denote the protocol type: heralding on two photons as in this work, on a single photon, direct state transfer though laser-mediated cavity interactions, and one instance of a simple distillation protocol applied to single-photon heralds (in conditions of "extreme photon loss" [CB08]). See table 6.4 for references.

	Fidelity	Rate	Wavelength	
Maryland (2007)	0.63(3)	$0.002\,\mathrm{s}^{-1}$	369 nm	ref. [Moe+07]
Maryland (2009)	0.87(2)	$0.0015\,\mathrm{s}^{-1}$	369 nm	ref. [Mau+09]
München (2011)	0.95(3)	$0.0024\,\mathrm{s}^{-1}$	795 nm	ref. [Let+11]
München (2012)	0.85(1)	$30\,\mathrm{s}^{-1}$	795 nm	ref. [Rit+12]
München (2012)	0.781	$0.027\,\mathrm{s}^{-1}$	795 nm	ref. [Hof+12]
Innsbruck (2013)	0.64(2)	$0.23\,\mathrm{s}^{-1}$	493 nm	ref. [Slo+13]
München (2013)	0.79(1)	$10\,\mathrm{s}^{-1}$	780 nm	ref. [Nöl+13]
Delft (2013)	0.69(5)	$0.0017\,\mathrm{s}^{-1}$	637 nm	ref. [Ber+13]
Delft (2014)	0.87	$0.004\,\mathrm{s}^{-1}$	637 nm	ref. [Pfa+14]
Maryland (2014)	0.78(3)	$4.5\,\mathrm{s}^{-1}$	369 nm	ref. [Huc+15]
Delft (2015)	0.92(3)	$0.0003\,\mathrm{s}^{-1}$	637 nm	ref. [Hen+15]
Zürich (2015)	0.55(4)	$2300\,\mathrm{s}^{-1}$	950 nm	ref. [Del+16]
Delft (2017)	0.55(1)	$0.054\,\mathrm{s}^{-1}$	637 nm	ref. [Kal+17]
Cambridge (2017)	0.62(2)	$7300\,\mathrm{s}^{-1}$	968 nm	ref. [Sto+17]
Delft (2018)	0.81(2)	$6\,\mathrm{s}^{-1}$	637 nm	ref. [Hum+18]
Oxford (2020)	0.940(5)	$180\,\mathrm{s}^{-1}$	422 nm	ref. [Ste+20]
Delft (2021)	0.816(4)	$9\,\mathrm{s}^{-1}$	637 nm	ref. [Pom+21]
Oxford (2021)	0.960(1)	$100\,\mathrm{s}^{-1}$	422 nm	ref. [Nad+22]
München (2021)	0.89(2)	$0.012\,\mathrm{s}^{-1}$	780 nm	ref. [Zha+22]
München (2022)	0.82(2)	$0.072\,\mathrm{s}^{-1}$	1517 nm	ref. [vLee+22]
Innsbruck (2022)	0.85(4)	$0.052\,\mathrm{s}^{-1}$	854 nm	ref. [Kru+23]

Table 6.4: Bell-state fidelity and rate for entanglement generated using an optical link between remote matter qubits (see fig. 6.16).

6.4.2   Other remote entanglement experiments

There are several possible dimensions along which remote qubit entanglement experiments can be evaluated. The most fundamental, and most widely reported, metrics are the quality of the generated entanglement, most commonly expressed in terms of the fidelity of the produced state to the closest maximally entangled state, which coincides with the fully entangled fraction $\mathcal{F}$, and the rate at which these entangled qubit pairs can be produced[10]. A survey of previous demonstrations of entanglement between remote mat-

---

[8]Cursory experiments attempting to disentangle this were hampered by the fact that at the time, both alignment and gate errors would drift at the sub-percent level.

[9]In fact, a calculation using the best observed ion–photon density matrices predicts a slightly *worse* ion–ion Bell state error, $\mathcal{E} = 4.4(2)\,\%$, than the best measured data actually shows, $\mathcal{E} = 4.0(1)\,\%$. Any attempt at interpretation of this barely significant difference is at best tenuous as the ion–photon and ion–ion data were not acquired contemporaneously and the timing of the experimental sequences differed slightly, but possible explanations for such a discrepancy would include detector jitter or common-mode noise between the two nodes such as on the phase of the 674 nm laser that would cancel in ion–ion experiments.

[10]Other quantities of interest are the ratio of coherence times available in the particular qubit system to the time necessary to generate one Bell pair (which limits the size of the computations/multi-qubit entangled

ter qubits mediated by an optical link along these two dimensions is shown in fig. 6.16, with extra numerical information in table 6.4. At the time of writing, the ion–ion entanglement results presented here constitute the highest-fidelity optical remote entanglement demonstration across all qubit platforms, and the highest generation entanglement rate for fidelities larger than 65 %.

In trapped ions, remote entanglement experiments were pioneered by the group of C. Monroe at the University of Maryland[11] [Moe+07]. Compared to their latest results [Huc+15], which were obtained in ^{171}Yb$^+$, our results represent a 5.5× reduction in error at a 22× increased rate. All these results (including ours) were obtained using blue/ultraviolet spontaneous decay photons for nodes located in the same laboratory. Very recently, results across 230 m distance (520 m fibre) were reported from the University of Innsbruck [Kru+23], for which cavities were used to enhance coupling to the 854 nm $P_{1/2} \leftrightarrow D_{3/2}$ transition in ^{40}Ca$^+$. Owing to limitations of the cavities, only much lower fidelities and rates of 0.85(4) and 0.052 s^{-1} were achieved. Downconversion to wavelengths more suitable to long-distance fibre transmission has been demonstrated in a number of separate experiments (see discussion in §8.1.2), but thus far not in an integrated two-node remote entanglement experiment.

Using optically trapped rubidium atoms, two groups in Munich (G. Rempe, H. Weinfurter) have demonstrated a wide variety of remote entanglement generation techniques, both using free-space collection of emission as well as optical cavities. The strong coupling between atoms and light achievable using cavities also enabled the demonstration of non-demolition detection of photons [Nie+21], which could be of wider interest for quantum networking applications. Recently, entanglement between two Rb atoms mediated by photons down-converted to 1517 nm was demonstrated [vLee+22], in the first such experiment for two matter qubits. The rates and fidelities achieved in neutral atom experiments have remained lower than demonstrated here [Rit+12; Zha+22].

Nitrogen-vacancy centres in diamond have also enjoyed popularity as a candidate qubit platform for quantum networks, and have prominently been investigated at Delft University of Technology (R. Hanson). While high fidelities have been reached (~ 90 %) using two-photon herald methods, this was at extremely low rates ($\sim$ mHz) [Pfa+14; Hen+15], as the small branching fraction of ≈ 3 % for decays that do not also involve vibronic bulk interactions (the "zero-phonon line") limits the achievable collection efficiencies [Dav74; Ruf+21]. Most recent work in NV centres has thus explored single-photon heralding schemes, which require global optical phase stability across the network link. Thus far, the fidelities achieved this way were limited to ≈ 80 % (e.g. in ref. [Pom+21], a demonstration of a three-node network, for which six cascaded optical phase stabilisation loops were necessary).

---

states accessible in the system), and, for quantum communication applications, the physical distance between the nodes.

[11] The group recently relocated to Duke University (Durham, NC, United States).

In solid-state quantum dots, remote entanglement generation rates of several kilo-hertz were demonstrated [Del+16; Sto+17], at present the fastest rates across all optical remote entanglement experiments[12]. The short dephasing times and other technical factors have thus far limited state fidelities to barely above the classical limit of fifty percent, however. An interesting technical facet of these quantum dot experiments is the use of electro-optical modulators to modulate the necessary initialisation and excitation lasers to achieve repetition rates in excess of 10 MHz, which could potentially be of interest for experiments in trapped ions as well.

In conclusion, informed by an in-depth theoretical analysis (see chapter 4), and by careful engineering of the experimental implementation, we were able to significantly improve on the state of the art in high-rate, high-fidelity remote entanglement generation. Crucially, the performance routinely achieved in the current configuration of the experiment is already sufficient to enable interactive exploration of the entanglement generation process[13], as well as to demonstrate rudimentary quantum networking applications in a realistic setting. In the next chapter, we will turn our attention to just such an application – device-independent quantum key distribution, which requires the production of millions of Bell pairs at high fidelity, and thus also stable operation of the remote entanglement link –, before finally discussing further potential improvements to rate and fidelity in chapter 8.

---

[12] Entanglement rates in excess of 10 kHz were achieved between superconducting qubits, but through the use of microwave photons, requiring the use of waveguides cooled to mK temperatures.

[13] From an experimentalist's perspective, a repetition rate of $100 \, \text{s}^{-1}$ for simple experiments requiring only one Bell pair is already twice as fast as many sensitive experiments, often synchronised to the AC mains cycle, would run.

7 *Device-independent quantum key distribution*

We use the remote entanglement link described in the preceding chapters to implement a long-sought-after form of quantum cryptography that approaches the ultimate limit on security by avoiding any dependency on the detailed workings of the quantum devices involved. This task is known as device-independent quantum key distribution, short DIQKD. Our experiment is the first successful realisation of such a protocol, demonstrating that it is feasible for two parties connected only by a (quantum-) optical fibre link to obtain a private encryption key in a device-independent fashion, even given the imperfections of present-day technology. In doing this, we rely on both the high-fidelity, high-rate entanglement and the associated large detection-loophole-free Bell violation accessible in our experiment, but also on significant improvements to the theoretical analysis and classical algorithms required for the parties to obtain a fully secure, identical key from the raw experimental data.

This chapter is intended as a fairly self-contained introduction to entanglement-based quantum key distribution accessible to the experimental quantum physicist, while also providing enough details on the practical implementation to substantiate the claim of a first real-world demonstration. After a brief high-level introduction to quantum cryptography and the role of entanglement (§7.1), we more formally consider the device-independent setting, from its remaining assumptions to the precise notion of security (§7.2). A concrete DIQKD protocol is described; we discuss elements of its security analysis to illustrate the challenging requirements imposed on experimental platforms (§7.3). The remainder of the chapter focusses on the particular implementation, including enough detail on both quantum and classical components for the results to be replicated on other remote entanglement platforms in the future (sections 7.4–7.6), before briefly discussing concurrent and future work as well as the implications for real-world security (§7.7).
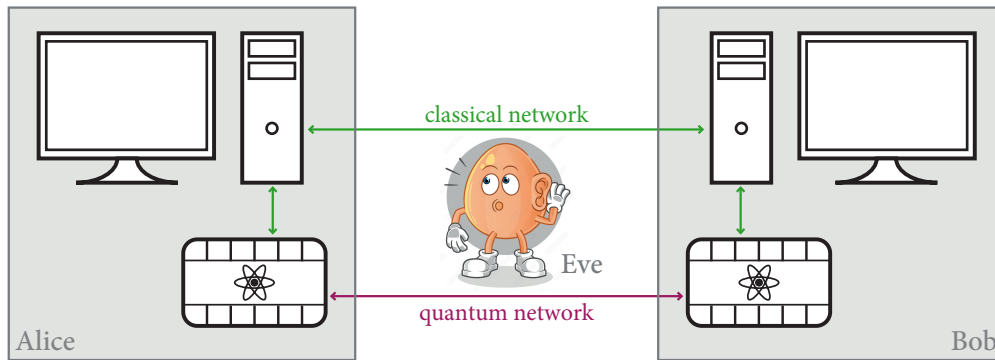
Figure 7.1: Cartoon illustration of the two-party encryption scenario. Two parties, Alice and Bob, wish to privately exchange messages over a public classical network (e.g. the internet). However, an arbitrarily powerful eavesdropper, Eve, has access to all information beyond the boundaries of Alice's and Bob's secure sites, which are indicated by the boxes. If Alice and Bob each have access to one of a pair of quantum devices connected via a (public) quantum channel, quantum key distribution (QKD) can achieve information-theoretic security.

7.1 A conceptual introduction to entanglement-based quantum cryptography

Perhaps the most fundamental scenario in cryptography is that of two-party secure communication. Two distant parties – we will call them Alice and Bob, following the time-honoured conventions of the field – wish to exchange information over a public (untrusted) network connecting their locations. This ability to privately communicate using shared network infrastructure is of fundamental importance in the modern information age, which has fuelled considerable innovation in the age-old struggle between "code makers" (cryptographers) and "code breakers" (cryptanalysts). Its culmination are a number of classical key exchange protocols [DH76; RSA78; Mer78], which are in ubiquitous use today to allow two remote parties to establish encryption keys for secret communication, for example all across the internet.

These protocols, however, rely on unproven conjectures regarding the computational hardness of certain problems, for instance, that any eavesdropper, which we shall refer to as Eve, cannot feasibly factor a large number into its prime factors. If Eve's computational capabilities are unbounded, the only way to achieve security is for Alice and Bob to use a pre-shared encryption key that is exactly as long as the messages, and only use that key once [KL15]. In this case, the *one-time pad* technique, illustrated in fig. 7.2, achieves perfect information-theoretical security, but the challenge of distributing the key material ahead of time makes it impractical for many situations.

Nevertheless, the one-time-pad is enormously useful, as it allows the parties to separate the message encryption in time from the key exchange. If Alice and Bob have the opportunity to exchange a long key at some point, for example through physical proximity in the real world (or, in a spy movie, an armed courier with a suitcase handcuffed to their wrist),

Encryption

message **m**	0 1 0 0 1 1 1 0 1 1 1 1 0 1 0 1 0 1 1 0
key **k**	⊕ 1 0 1 1 1 0 1 0 1 0 1 1 1 1 0 0 1 1 1 1
ciphertext **c**	1 1 1 1 0 1 0 0 0 1 0 0 1 0 0 1 1 0 0 1

Decryption

ciphertext **c**	1 1 1 1 0 1 0 0 0 1 0 0 1 0 0 1 1 0 0 1
key **k'**	⊕ 1 0 1 1 1 0 1 0 1 0 1 1 1 1 0 0 1 1 1 1
message **m'**	0 1 0 0 1 1 1 0 1 1 1 1 0 1 0 1 0 1 1 0

Figure 7.2: One-time pad encryption for information-theoretic security. To encrypt a message m, Alice computes the bitwise exclusive-or $c = m \oplus k$ with a shared key k. For decryption, the procedure is simply repeated, $m' = c \oplus k'$. If Bob has access to the same key, $k' = k$, he recovers the message $m' = m$. However, to an adversary not in possession of the key the ciphertext is completely random (and any arbitrary message m' of the same length would be a valid decryption for the key $k' = c \oplus m'$). Thus, the one-time pad is provably secure.
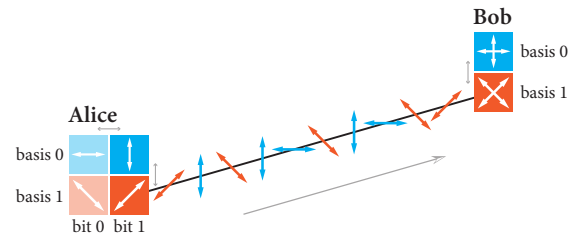
Figure 7.3: "Regular" prepare-and-measure QKD using polarisation encoding [BB84]. Alice encodes each bit of the random key in a single photon which she sends to Bob, choosing one of two polarisation bases at random. Bob similarly chooses a random measurement basis each time. Alice and Bob later compare their basis choices and discard mismatched rounds. Since Eve does not know the basis choices ahead of time, any measurements to eavesdrop on the transmitted photons would risk her altering their state, which Alice and Bob can detect. As Alice relies on her source emitting only a single photon each time (and in the correct basis), this protocol cannot offer any device-independent security. (Graphical presentation inspired by that in ref. [Sto11].)

they can, potentially much later, use it to promptly exchange messages. In other words, because the one-time pad is information-theoretically secure, a pre-shared key of length l and an authenticated (but insecure) channel for l bits can be used to *construct* a secure channel for l bits; this notion of construction can in fact be formalised in a mathematically precise way [Mau11].

The one-time pad key can—and must—be completely independent of any of the contents of later messages. This property enables what has become the best-known example of *quantum cryptography*, that of *quantum key distribution* (QKD): Quantum mechanics promises the ability for two parties to generate secret correlations without any third party learning of the results. If a secret key has been established in such a QKD process, the one-time pad then completes the link to *quantum-secure encryption*, i.e. an information-theoretically secure cryptosystem for the encryption of messages, where no[1] initial key is required, but generated on-the-fly from shared quantum resources instead.

One intuitive way to implement QKD is based on the existence of non-commuting measurements and the non-cloning property in quantum mechanics. The prototypical example for this is the original proposal by Bennett and Brassard [BB84], the principle of which is easily explained (see fig. 7.3): Alice sends single photons to Bob in one of two mutually unbiased polarisation bases. As she chooses the basis at random each time, and Bob does

---

[1] In fact, a short initial key is required for authentication purposes, as discussed in §7.2.1. A more precise name for quantum key distribution would have perhaps been "quantum key expansion".

the same with his measurement basis, Eve cannot measure the photons herself without affecting their polarisation state, leaving a statistical signature of her eavesdropping. Such protocols, with some adaptions to better match available technology (e.g. in the form of "decoy state" protocols that work with weak coherent laser pulses instead of single photons [LMC05]), have been widely studied and even deployed in commercially available systems; see ref. [Pir+20] for a recent review. However, such prepare-and-measure protocols are in their security analysis inevitably based on device models that describe the behaviour of the sources and detectors, and which need to be trusted in the security analysis. An accurate characterisation of the devices on the quantum level, valid over time, can be challenging to obtain, however, and any mismatch between the actual device behaviour and the assumptions made at the time of the security analysis can leave a loophole for Eve to gain information about the outcomes. A number of such attacks against QKD systems have in fact been demonstrated in practice [Zha+08; Lyd+10; Ger+11; Wei+11; GSM20], a discipline sometimes referred to as *quantum hacking*.

Curiously enough, it is possible to construct QKD protocols that do not require any detailed knowledge about the behaviour of the quantum devices that they are executed on. This, perhaps counter-intuitive, possibility of *device-independent security* builds on another facet of quantum mechanics: that of non-locality in the form of entanglement. Such protocols go back to an early proposal by A. Ekert [Eke91] suggesting the use of a Bell inequality violation between the outputs of Alice's and Bob's devices to bound the information available to any adversary. In the rest of this section, we will develop an intuitive perspective on how the observation of a particular kind of correlations in a Bell test experiment can conceptually lead to guarantees about the privacy of the measured outcomes, before returning to a more formal discussion in §7.2.

7.1.1   *Magic coins and free choice*

The mechanism by which J. S. Bell's eponymous inequality [Bel64] shows that quantum mechanics is incompatible with local realist hidden-variable theories is relatively straightforward and has been discussed at length in the literature [Gen05], as is their role in the recent experiments which finally demonstrated that Nature in fact does not behave in a locally realist way [Giu+15; Hen+15; Sha+15]. To motivate the role of Bell inequalities in the context of device-independence, however, it is instructive to consider a slightly different line of argumentation (as e.g. done in ref. [ER14]), which initially does not even make any references to quantum theory.

Suppose Alice and Bob each have a pair of special, "magic" coins, of which they are allowed to choose one, and only one, to throw. Let $A_i, B_i \in \{-1, 1\}$ be the outcome (-1 heads, 1 tails) observed when choosing to throw the i-th coin ($i \in \{0, 1\}$). The coins are magically linked such that, depending on which one of their pair each party chooses, their

outcomes are always correlated in the following way:

$$A_0 = B_0, \quad A_0 = B_1, \quad A_1 = B_0, \quad A_1 \neq B_1. \tag{7.1}$$

Remarkably, assuming that Alice and Bob really can choose between the coins freely and privately (without their choice leaking to a third party), these conditions alone already imply that there can be no third coin E that replicates the outcomes of one of the other coins. Consider, for instance, that a would-be eavesdropper Eve has a coin E that, by whatever other magical means, always produces the same outcomes as A_0. If now Eve throws her coin but Alice actually chooses to throw A_1, and for the sake of the argument, both coins come up as 1, then $A_0 = E = A_1$. But now, if Bob were to choose to throw B_1, then eq. (7.1) would result in the contradiction $B_1 \neq B_1$. Thus, the presence of Eve's copy E would force the choice of B_0 onto Bob for logical consistency. Similar contradictions ensue for other choices of coins and outcomes[2]. In other words, just the ability of Alice and Bob to choose their coins freely is enough to ensure that Eve cannot learn the outcomes of their throws, no matter how exactly the "magic" correlations actually come about.

While magic of this sort is unfortunately not real, quantum mechanics gets close: If we define a sum of expectation values mirroring eq. (7.1),

$$S := \langle A_0 B_0 \rangle + \langle A_0 B_1 \rangle + \langle A_1 B_0 \rangle - \langle A_1 B_1 \rangle, \tag{7.2}$$

the perfectly correlated coins would be equivalent to $S = 4$. The parameter S is of course the expectation value that appears in the Clauser–Horne–Shimony–Holt formulation of a Bell inequality, $|S| \leq 2$, often called the CHSH *score* [Cla+69], and we might hope that we can derive some of the same privacy guarantees from an observation of suitably high values of S. By choosing to implement the coins as a choice of two measurements on a maximally entangled state $|\Phi^+\rangle = \frac{1}{\sqrt{2}}(|\uparrow\uparrow\rangle + |\downarrow\downarrow\rangle)$, taking

$$A_0 = Z, \quad A_1 = X, \quad B_0 = \frac{1}{\sqrt{2}}(Z + X), \quad B_1 = \frac{1}{\sqrt{2}}(Z - X) \tag{7.3}$$

yields $S = \langle \Phi^+ | A_0 B_0 + A_0 B_1 + A_1 B_0 - A_1 B_1 | \Phi^+ \rangle = 2\sqrt{2}$. This evidently violates the CHSH inequality, and is in fact the maximum achievable score in quantum mechanics, known as Tsirelson's bound[3] [Cir80].

Before investigating whether there "is any magic left" in the quantum maximum, let us mention an equivalent way of phrasing this in terms of a two-player game, the CHSH *game*. Two parties, who cannot communicate, are each provided with one input x and y

---

[2] The conditions from eq. (7.1) nevertheless are not themselves contradictory, as without such an extra coin E, Alice and Bob only get to choose one of their coins each; A_0 and A_1 or B_0 and B_1 are never both "evaluated" simultaneously.

[3] The cyrillic name of the Russian–Israeli mathematician for which the $S \leq 2\sqrt{2}$ bound is named has been transliterated in two different ways; earlier work such as ref. [Cir80] uses "Cirel'son", while the author later himself preferred "Tsirelson" in his English-language publications.
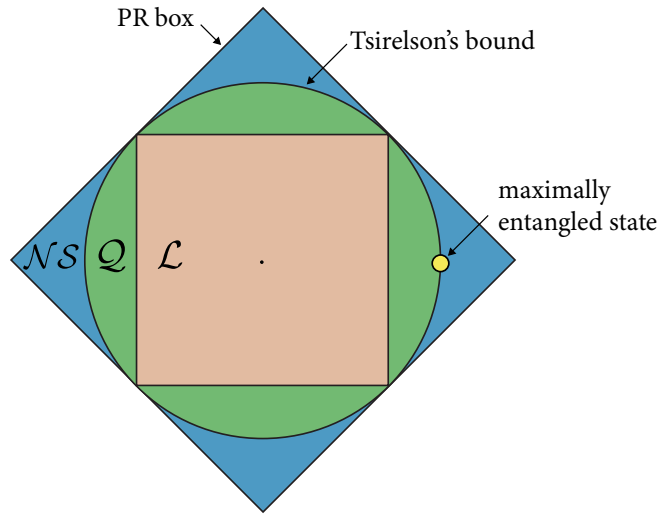
Figure 7.4: Conceptual illustration of the space of correlations allowed by non-signalling ($\mathcal{NS}$), quantum ($\mathcal{Q}$) and classical (realist local, $\mathcal{L}$) theories. The outer border corresponds to the "Popescu–Rohrlich box" [PR94] which always wins the CHSH game; the dot marks an extremal point in the space of quantum correlations, realised by a measurement on a maximally entangled state. (This figure was adapted from ref. [Bru+14], where the two coordinate axes are given a quantitative interpretation as the CHSH scores for two complementary Bell inequalities.)

randomly chosen from $\{0, 1\}$ and asked to produce outputs $a, b \in \{0, 1\}$ aiming to ensure that the relation

$$a \oplus b = x \cdot y \tag{7.4}$$

always holds (where $\oplus$ is modular addition in $\mathbb{Z}/2\mathbb{Z}$). This game can always be won using the magic coins from eq. (7.1), but the optimal quantum strategy is again to obtain a and b from a joint measurement $A_x \otimes B_y$ on the maximally entangled state $|\Phi^+\rangle$ (or, more precisely, $\frac{1}{2}(A_x + 1) \otimes \frac{1}{2}(B_y + 1)$ to convert the -1 eigenvalue to the 0 outcome). The probability ω of winning this game is then given by

$$\omega = \frac{S + 4}{8}. \tag{7.5}$$

7.1.2 Geometry of correlations

Even though a statistical test of the CHSH score S is at the heart of both "Bell test" experiments to reject a realist local description of Nature and the statement about the secrecy of outcomes that is of interest here, the role of the parameter is subtly different. This is most easily illustrated by considering the space of correlations – in our case, expressed by the conditional probabilities $P(a, b|x, y)$ – allowed by different physical theories, as done in fig. 7.4. The no-signalling condition $\mathcal{NS}$, forbidding the superluminal transfer of information, places a very general restriction on the possible correlations between separate systems. Much more restrictive than that are those imposed by classical physics $\mathcal{L}$, that

is the insistence on a realist local description, for which $|S| \leq 2$. The set of correlations allowed by quantum theory, $\mathcal{Q}$, is larger than that of classical physics, but is still more restricted than that of all non-signalling correlations.

Within quantum theory, correlations that are "maximally quantum" and saturate Tsirelson's bound $|S| \leq 2\sqrt{2}$ have an interesting property, that of *monogamy*: If two systems share such maximal correlations, they cannot be at all correlated with any third system [Bru+14]. In particular, this holds for any system that would be controlled by Eve, so that we have located some residue of the above "magic" in the quantum case. This can also be formulated in terms of the quantum states underlying the measurement results as *monogamy of entanglement* [CKW00] (and generalises to non-locality in a wider sense [Ton09]). Importantly for our case, if Alice and Bob observe a CHSH score of *almost* $2\sqrt{2}$, then Eve can have *almost* no information on their outcomes; the property degrades in a gradual way (see the later fig. 7.7).

In summary, in a Bell test experiment, observing a CHSH score $|S| > 2$ allows one to conclude that correlations in Nature are not confined to $\mathcal{L}$. For DIQKD, however, we assume that Nature is in particular described by quantum theory, such that observing $|S| \approx 2\sqrt{2}$ allows Alice and Bob to conclude that their devices operate close to the boundary of $\mathcal{Q}$, where, as described by the monogamy of entanglement, their outputs are hidden from Eve's prying eyes.

7.2 *Device-independent security*

To transform this theoretical insight about the nature of quantum correlations into a piece of quantum technology, Alice and Bob in their secure sites each operate one of a pair of quantum devices. The devices will internally make use of entanglement in some way, but the "user interface" Alice and Bob have to these devices is strictly classical: as illustrated in fig. 7.5, device-independent key distribution protocols proceed in rounds, where Alice and Bob each time present the device with a classical input, and get back one of two classical outputs. The expectation is that the inputs will affect the choice of measurements on a maximally entangled state shared between the two devices in a way that maximises the CHSH score, but this does not need to be an assumption: if the outputs follow the desired distribution, then it is clear from quantum theory that they can only result from measurements on a maximally entangled state, guaranteeing security from the monogamy of entanglement.

This concept of *self-testing* was first explicitly formalised in ref. [MY04], and remains a useful tool for conceptualising these experiments (although we do not take their mathematical perspective in the actual security analysis). Since Alice and Bob are ultimately not interested in certifying that their devices produce a maximally entangled state, but wish to obtain a shared, secure encryption key, a purely practical improvement is for them to most of the time choose inputs that yield perfectly correlated outputs, and only allocate
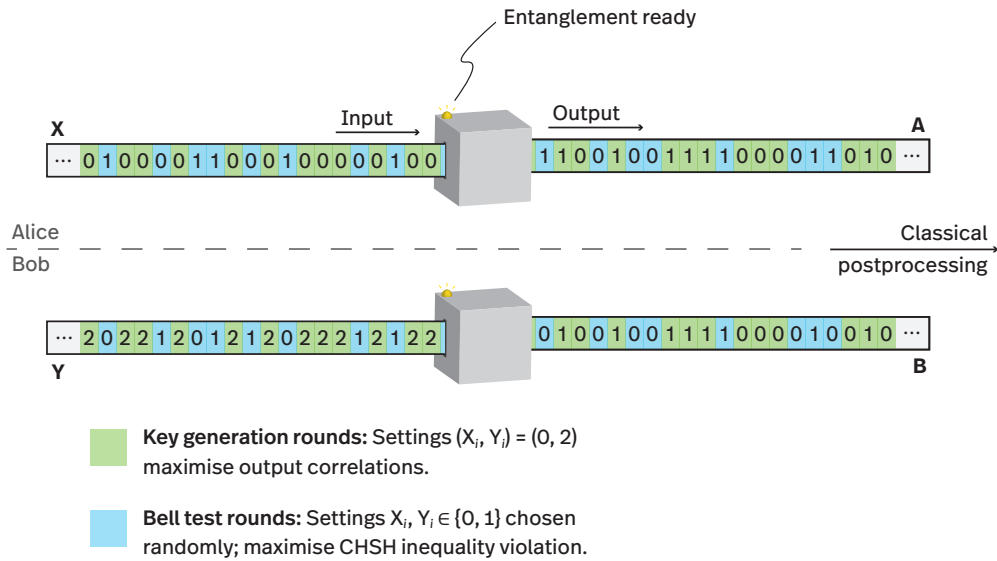
Figure 7.5: A high-level view of device-independent quantum key distribution. Alice and Bob each have a box which promises to generate and measure a maximally entangled state. They proceed in rounds; each time entanglement is available, they apply classical inputs X_i and Y_i to their respective box, and record the corresponding classical outputs A_i and B_i; no further assumptions about the internal workings of the quantum devices are made in the security proof. Rather, by randomly alternating between measurement settings realising a Bell test and settings which lead to highly correlated outputs, the output bit strings can be certified to originate from appropriate measurements on a quantum state that is close to being maximally entangled, ensuring secrecy.

some of the rounds to Bell tests (indicated by the colours in fig. 7.5). This way, they can generate key material much more efficiently than through the imperfect CHSH correlations while, crucially, still reaching statistical confidence that their devices operated on a maximally entangled state for all the rounds, as they choose between round types at random every time. The vital ingredient that allows Alice and Bob to be confident that such a procedure, and indeed also the Bell test approach itself, is secure – that is, the "root of trust" for DIQKD in general – is the fact that they make these choices based on trusted random number generators.

Continuing the intellectual lineage of Ekert's original proposal for entanglement-based QKD [Eke91], a number of concrete scenarios and protocols along these lines have been discussed in recent years, see e.g. refs. [Ací+07; VV14; ARV19] and references therein. The primary challenge lies in devising protocols and security proofs that do not make more *assumptions* than necessary, yet are able to guarantee *security* even with real-world, noisy quantum devices. Before turning to a description of the concrete protocol that we use to bridge the gap between assumptions and security in our demonstration, we first inspect these two, starting point and goal, in more explicit detail.

7.2.1 *Assumptions*

In our security analysis, we want to confine ourselves to the minimal set of assumptions usually employed in a DIQKD setting, as e.g. discussed in ref. [Pir+09]:

  (i) Any attacker is restricted to quantum physics.
 (ii) Alice's and Bob's laboratories are isolated from the outside world, that is, no unwanted information can leak to the environment,
(iii) they each have trusted classical computers for storage/post-processing,
 (iv) they share an authenticated (not necessarily secure) classical channel, and
  (v) they each have a trusted random number generator.

Conventional QKD adds an additional assumption [Sca+09; LCT14]:

 (vi) Alice's and Bob's quantum devices perform the expected operations, that is, they know an exact description of the quantum states and measurements.

Assumption (vi) could also be called that of *trusted device models*, and e.g. also includes the dimensionality of the involved quantum states [AGM06]. As discussed in the introduction to this chapter, this is not easy to achieve in practice, and it is precisely this (and only this) assumption that DIQKD avoids. So-called measurement-device-independent protocols sit somewhere in the middle; they allow untrusted measurements and thus, for instance, tolerate attacks on single-photon detectors used, but still require well-characterised sources.

Most physicists would be content with assumption (i), quantum theory being as well-tested as it is, and with this assumption only concerning the structure of reality in general, not a description of the particular quantum devices used. Nevertheless, it is interesting to note that there are potential ways of relaxing this assumption even further, e.g. by only requiring a more general no-signalling condition as discussed in [BHK05].

Assumptions (ii) and (iii) are intrinsic to the setting of cryptography. While appropriate isolation (including all possible side channels such as the timing of external communication) can be challenging to ensure in practice, even just the definitions of concepts such as "private keys" and "secrecy" implicitly assume that private information can be kept privately.

Assumption (iv) is not always explicitly stated, as it is fundamentally necessary for security in an elementary way: if Alice and Bob do not have a way of identifying each other, a trivial man-in-the-middle attack is possible, where Eve just honestly executes the DIQKD protocol with either of them, pretending to be the respective other party. Fortunately, information-theoretically secure authentication can easily be achieved even through classical means: while a short secret shared by Alice and Bob ahead of time is required (and consumed by the protocol), its length is only required to grow logarithmically in the size of the messages to secure [KL15]. In the DIQKD protocol described later in this chapter, we explicitly include classical signatures for all the relevant messages in the DIQKD protocol

(and budget for the corresponding secret key consumption)[4], just implicitly assuming the existence of a public (classical) channel.

Assumption (v) is critical, as the free choice of classical inputs underpins all the device-independent security claims. It is universally accepted that such a truly random choice can be implemented by a self-contained module that internally exploits some physical process that is to the best of our knowledge random, e.g. by measuring the photon counting statistics of a weak beam at a beamsplitter. To avoid the dependence on a particular device implementation, one could employ device-independent randomness extraction techniques [KA20] to increase trust in the random sources, or even cleverly re-use other randomness unknown to Eve [Agu+16]; this does, however, not change the fact that "free choice" is ultimately necessary for all techniques based on Bell inequalities.

7.2.2   Security definitions: soundness and completeness

A DIQKD protocol necessarily has two outcomes: either the protocol succeeds and Alice and Bob have generated a private key, or it aborts, for instance because Eve has interfered with the execution. To formalise the desired outcome, we denote the successful completion of the protocol as Ω, and write $\rho_{|\Omega}$ for the (normalised) quantum state of the entire scenario, comprising Alice's and Bob's private subsystems A and B (their nodes), as well as those parts available to a potential adversary E.

Any useful security definition for a DIQKD protocol is going to be a probabilistic statement. This might be unavoidable due to the probabilistic nature of quantum measurements or noise in the experimental realisation, but more fundamentally, is a direct consequence of the statistical nature of the self-testing approach: The number of Bell test rounds in an execution of the protocol places an upper bound of the amount of deviation from an honest implementation we can hope to detect. We thus introduce the notions of *correctness*, *secrecy*, *soundness* and *completeness* each "up to a small deviation" given by a parameter $\epsilon \in [0, 1]$, following e.g. ref. [ARV19]. Sensible protocols would allow these parameters to be chosen to be arbitrarily small, although at increased cost e.g. in terms of minimum number of rounds required.

If the protocol succeeds, the private keys K_A and $K_B \in \{0, 1\}^l$ that Alice and Bob end up with should agree. We formalise this in the notion of ϵ_{corr}-*correctness*, for which we require

$$P(\Omega \wedge K_A \neq K_B) \leq \epsilon_{\text{corr}}. \tag{7.6}$$

This shared key should be completely unknown to the adversary. To capture this, we define a protocol to be ϵ_{sec}-*secret* if

$$P(\Omega) \frac{1}{2} \|(\rho_{|\Omega})_{K_A E} - \mathbb{U}_{K_A} \otimes (\rho_{|\Omega})_E\|_1 \leq \epsilon_{\text{sec}}, \tag{7.7}$$

---

[4]This is the reason why we do not explicitly mention authentication when introducing the other requirements in the journal manuscript, ref. [Nad+22].

that is, the joint state of adversary and secret key $(\rho_{|\Omega})_{K_A E}$ is very close (in trace distance) to the state where K_A just factors out as maximally mixed, i.e. a uniform distribution $\mathbb{U}_{K_A}$ over all possible strings. ϵ_{sec}-secrecy necessarily implies that the produced key is (almost) uniformly random.

The desired success criterion is then the logical conjunction of the two: Alice and Bob share the same, random key, that is unknown to Eve. We call a protocol that fulfils both criteria ϵ_{snd}-*sound*, requiring

$$\epsilon_{\text{corr}} + \epsilon_{\text{sec}} \le \epsilon_{\text{snd}}. \tag{7.8}$$

Going only by the above criteria, one protocol that would be ϵ_{snd}-sound to an arbitrary degree and for arbitrary devices would be the trivial protocol that always aborts – perfectly secure by never generating a key. We hence define our last desired property: A protocol is ϵ_{com}-*complete* if it admits a honest implementation which aborts with probability less than ϵ_{com}, that is

$$P(\Omega) \ge 1 - \epsilon_{\text{com}}. \tag{7.9}$$

To meet these conditions to a meaningful extent when implemented using real-world, noisy devices will require some amount of error correction for correctness, and the ability to tolerate quite substantially non-ideal CHSH scores while preserving secrecy. This is achieved by combining the "quantum" data acquisition phase with classical post-processing methods amenable to an information-theoretic treatment. Such a combined protocol is presented in the next section.

7.3   The concrete protocol

Our concrete DIQKD protocol follows the same overall structure as in a number of previous proposals (see e.g. refs. [Pir+09; Arn+18]). The complete structure, including the classical post-processing steps, is illustrated in fig. 7.6.

Before execution starts, all the relevant protocol parameters, n, γ, m and ω_{thr}, are fixed, which we shall define shortly. Alice and Bob also need to share a short private key K_0 for authentication purposes, as we do not make the prior assumption of an authenticated channel (see §7.2.1). Its exact length depends on the implementation details of the classical post-processing pipeline.

The data acquisition step proceeds in one block, consisting of n rounds, that is, n uses of the quantum "boxes". For each round, Bob randomly decides whether to perform a Bell test (with probability given by the test round fraction $\gamma \in (0, 1)$) or a key generation round (with probability $1 - \gamma$). After the quantum devices indicate that entanglement has been established, and have been disconnected from the environment, Bob communicates the choice of round type in form of the test round bit T_i to Alice. This avoids later "sifting", where some rounds without useful correlations due to a mismatch in measurement bases

Protocol parameters		Fix protocol parameters n, γ, m, ω_{thr} before starting.
Pre-shared key	K_0	Alice and Bob share a key K_0.
Private randomness	X Y	Alice holds $\mathbf{X} = (X_i)_i$ sampled according to $P(X_i = 0) = P(X_i = 1) = 1/2$. Bob holds $\mathbf{Y} = (Y_i)_i$ with $P(Y_i = 0) = P(Y_i = 1) = \gamma/2$ and $P(Y_i = 2) = 1 - \gamma$.
Data acquisition	$\times n$ T_i X_i Y_i A_i B_i	For every round $i \in \{1, \ldots, n\}$: » Alice and Bob wait for the heralding signal, then disconnect the link. » Bob computes $T_i = \chi(Y_i \neq 2)$ and sends T_i to Alice. If $T_i = 0$, Alice overrides X_i and sets it to 0. » Alice and Bob apply settings X_i and Y_i and record the measurement results as A_i and B_i.
Basis revelation	X	Alice sends $\mathbf{X}$ to Bob.
Error correction	M	Alice computes syndrome $\mathbf{M}$ of $\mathbf{A}$ (length m) and sends it to Bob. Bob reconstructs a guess $\tilde{\mathbf{A}}$ of $\mathbf{A}$.
Verify error correction	G_{EC} ✓	Alice computes an encrypted hash $\mathbf{G}_{\mathrm{EC}}$ of $\mathbf{A}$ using K_0 and sends it to Bob. Bob computes the hash $\tilde{\mathbf{G}}_{\mathrm{EC}}$ of $\tilde{\mathbf{A}}$. If $\tilde{\mathbf{G}}_{\mathrm{EC}} \neq \mathbf{G}_{\mathrm{EC}}$, the protocol aborts.
Verify Bell violation	✓	If Bob, considering only rounds i with $T_i = 1$, finds that $\sum_i \chi(A_i \oplus B_i \neq X_i \cdot Y_i) > n\,\gamma\,(1 - \omega_{\mathrm{thr}})$, the protocol aborts.
Privacy amplification	⊻ ⊻	Alice and Bob apply a strong randomness extractor with a seed from K_0 to $\mathbf{A}$ and $\tilde{\mathbf{A}}$ to obtain K_A and K_B.
Authentication	✓ G_B C G_A ✓	Bob computes an encrypted tag $\mathbf{G}_B$ for $\mathbf{T}$ using K_0 and an unconditionally secure message authentication code and sends it to Alice. Alice verifies the tag and sends $C = 1$ to Bob if it matches the earlier communication, $C = 0$ if it does not. Alice similarly computes a tag $\mathbf{G}_A$ for $(\mathbf{X}, \mathbf{G}_{\mathrm{EC}}, C)$ and sends it to Bob. Bob defines $F = 1$ if the tag matches and $C = 1$, $F = 0$ otherwise.
Key activation	F G_F ✓	Bob sends F to Alice. Bob computes a tag $\mathbf{G}_F$ for F using K_0 and sends it to Alice. Alice verifies the tag; if it does not match or $F = 0$, the protocol aborts.
Key concatenation	🔑 🔑	Alice removes the used bits from K_0 and defines K'_A by appending K_A. Bob does the same to get K'_B.
	Alice Bob	

Figure 7.6: Structure of our DIQKD protocol. Before execution starts, the number of rounds n, the probability γ of each round being chosen to be Bell test round, an acceptance threshold ω_{thr} for the CHSH winning probability, and the length m of the error correction syndrome are fixed. An initial key K_0, mostly reusable, is required to seed the privacy amplification and authentication algorithms, and as a one-time pad to encrypt a few short messages (indicated using a key symbol). Arrows indicate the classical messages exchanged between the parties, bold letters strings consisting of multiple bits, χ the indicator function with $\chi(P) = 1$ if P is true and 0 otherwise. In a honest implementation using noisy quantum devices, the Bell violation verification step is the last point of failure: if the classical communication between the nodes is not tampered with, the later steps deterministically succeed.

would be removed. For Bell test rounds, Alice and Bob each choose their inputs X_i and Y_i from $\{0, 1\}$ uniformly at random, for key generation rounds, they apply the fixed inputs $(X_i, Y_i) = (0, 2)$. In either case, they keep track of their outputs in bit strings $A = (A_i)_i$ and $B = (B_i)_i \in \{0, 1\}^n$.

The protocol then enters the purely classical post-processing or reconciliation phase, where Alice and Bob need to verify the amount of Bell violation, and extract an identical key. To achieve both those steps in the one-way reconciliation scheme employed here, one node (here: Bob) needs to know the outcomes of the other node (here: Alice). To achieve this with a minimum of information leakage, Alice sends a short error correction syndrome string M of length $m \in \mathbb{N}$ (along with her basis choices X), which allows Bob to reconstruct a guess at her outcomes $\tilde{A}$. To verify whether this was successful, Alice also sends an encrypted hash of A to Bob; if the hashes match, then $A = \tilde{A}$ with very high probability, if not, the security may be compromised and the protocol aborts. Now having $XYAB$, Bob can verify whether the CHSH winning probability observed exceeds a pre-agreed threshold $\omega_{\text{thr}} \in (0, 1)$ by simply tallying up the number of CHSH losses, that is, the test rounds where $A_i \oplus B_i \neq X_i \cdot Y_i$. If the number of losses exceeds the threshold, the protocol again aborts.

At this point, the parties know that – assuming their classical communication was not tampered with – their quantum devices behaved close enough to an ideal, honest implementation and they can locally process their strings A and $\tilde{A}$ into shorter strings K_A and K_B using a randomness extractor to obtain identical final keys with a guaranteed, arbitrarily low bound on the information available to any adversary. The amount of private entropy in the outcomes, that is, the length of the secret key that can be obtained in this privacy amplification process for a given soundness error ϵ_{snd}, only depends on the initially fixed parameters n, γ, ω_{thr} and m. This fixed structure makes it straightforward to formulate the statistical bounds in a sound way, and makes it relatively easy to disallow many adaptive strategies that the attacker might otherwise have.

Alice and Bob are then left with the task to ensure that their classical communication was not tampered with, which would mean that any of their earlier conclusions about secrecy and correctness might be faulty. To this end, each side simply computes an encrypted hash of all the previously sent communication and sends it to the other party to verify against their local records.

Finally, there is one wrinkle regarding the establishment of success: In the particular arrangement shown in fig. 7.6, Bob already knows at the end of the authentication step that his key is secure, and can thus be used in further communication, e.g. as a one-time pad. Alice, however, has no way of knowing whether the check of her encrypted signature actually succeeded on Bob's side, so Bob sends one last authenticated message to Alice to inform her whether the protocol aborted or not, at which point the QKD procedure is

complete[5], and both parties can be confident that the generated key is secure to use in other cryptographic applications.

7.3.1   Outline of the security proof

To prove that a particular instantiation of this protocol is secure, that is, that for a certain choice of parameters n, γ, ω_{thr} and m an ϵ_{snd}-sound key of length l can be extracted, is a non-trivial, highly technical affair. The interested reader might want to refer to the supplement to ref. [Nad+22] for the detailed analysis of the scenario in the later demonstration experiment, which has considerable overlap with recent, purely theoretical treatments [ARV19; Tan+20]. Here, we just outline enough of the details to discuss the consequences for the requirements on an experimental realisation.

The first of the two requirements for soundness, ϵ_{corr}-correctness per (7.6), is easy to achieve for arbitrary $\epsilon_{\text{corr}} > 0$ by an appropriate choice of hash function in the error correction verification step: if the hashes for A and $\tilde{A}$ match, then so with high probability do the strings themselves, and thus the keys K_A and K_B derived from them via privacy amplification. One such choice (universal hashing) is discussed in §7.4.3.

The crux of the issue lies in the second component of soundness, ϵ_{sec}-secrecy. Here, the independence of the key K_A from any side-information available to the eavesdropper demanded by eq. (7.7) is guaranteed by the randomness extractor in the privacy amplification stage. For a suitable choice of extractor, this can be achieved for extracted key lengths l that are given (up to a logarithmic correction) by $H_{\text{min}}^{\epsilon_s}(A|\overline{E})$, that is, the amount of conditional entropy[6] left in Alice's outputs considering all the information $\overline{E}$ available to Eve.

This "side-information" $\overline{E}$ consists of all the classical messages exchanged by Alice and Bob during the protocol, plus any quantum state E that Eve herself might hold. In terms of the symbols from fig. 7.6, we can thus write (substituting the more information-rich Y for the round-type messages T)

$$\overline{E} = X\,Y\,\underbrace{M\,G_{EC}\,G_B\,C\,G_A\,F\,G_F}_{=:M'}\,E.$$

In this information-theoretic setting, we assume that each bit of the classical messages leaks the maximum possible entropy of one bit of information about A to Eve, and can

---

[5] Although trivial, this final "key activation" procedure was, to my knowledge, not explicitly discussed in earlier work. The need to address what is otherwise an imbalance between the state of Alice and Bob emerged only during the practical implementation of the protocol, and because of the choice not to just assume an authenticated channel. By structuring the protocol such that Bob is always the one to decide to abort, he can simply encode the overall protocol success into the final key activation message. This requires an extra authenticated message, but leads to a pleasantly symmetric, composable postcondition for the protocol, where both parties can go on to use the generated key material.

[6] For technical reasons – to accurately describe the worst-case distinguishability of states in the finite-size, "one-shot" regime – the entropy measure that appears here is not the Shannon/von-Neumann entropy H, but the ϵ_s-smooth min-entropy $H_{\text{min}}^{\epsilon_s}$. The basic intuition for the quantities remains the same.
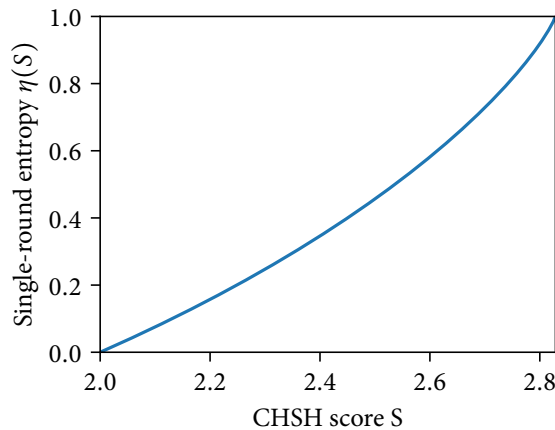
Figure 7.7: Amount of secret randomness $\eta(S)$ generated in a single round of the CHSH game, as quantified by the conditional entropy $H(A|XYE)$. $\eta(S)$ can be seen as the infimum across all device realisations compatible with an observed CHSH score S, describing the von Neumann entropy that can be guaranteed in the worst case. Some secret entropy is generated as soon as S exceeds the classical bound, $S > 2$; the amount gradually increases until the outcomes are known to be perfectly secret at $S = 2\sqrt{2}$.

thus estimate

$$H^{\epsilon_s}_{\min}(A|\overline{E}) \geq H^{\epsilon_s}_{\min}(A|XYE) - |M'|,$$

where $|M'|$ is the length of the exchanged messages, excluding those describing the classical inputs. The length of the authentication-related messages $(G_{EC}, G_B, C, G_A, F, G_F)$ is some constant k that depends on the hash function chosen (in our demonstration, we use 64-bit hashes, so $k = 258$). The length of the error correction syndrome M required depends on the efficiency of the chosen error-correction code, but will always be lower-bounded by the information-theoretic uncertainty that Bob has about Alice's outcomes. Considering, for the purposes of expositional clarity, an ideal error correction code and the asymptotic limit, we can thus bound the length of the extractable key as

$$l \approx H(A|XYE) - H(A|XYB) - k. \tag{7.10}$$

In other words, the key length emerges as a kind of *balance of entropies*, the difference between the uncertainty that Eve has about Alice's outcomes and the uncertainty that Bob has about the same (and the fixed protocol overhead).

To obtain a rigorous bound for the first term, $H^{\epsilon_s}_{\min}(A|XYE)$, is then the main focus of the security analysis, and, since we want to make minimal assumptions on the behaviour of the system, technically quite challenging. Herein lies a big division between various DIQKD protocols and security proofs. One hugely simplifying assumption is to assume that the entire data acquisition phase of n rounds can be described as n independent copies of a single round, where both device and eavesdropper behaviour (with associated quantum side information E_i) can be described as independent and identically distributed (i.i.d.). In this case, also referred to as that of *collective* attacks, the entropy can, up to a $\sqrt{n}$ statistical correction, easily be bounded by $n H(A_i|X_iY_iE_i)$, that is, just n times the entropy content of a

single round that is compatible with the observed device behaviour. If Alice's and Bob's devices behave according to an expected CHSH score of S, the worst-case entropy compatible with this is given in terms of the binary entropy function h (eq. (2.3)) as [Pir+09]

$$\eta(S) = 1 - h\left(\frac{1 + \sqrt{\left(\frac{S}{2}\right)^2 - 1}}{2} \right), \tag{7.11}$$

which is illustrated in fig. 7.7.

Unfortunately, such a restriction to *i.i.d.* behaviour is not particularly well-motivated in the generic device-independent case. As A. Ekert comments [Eke18]:

> At the time, even the most fervent advocates of the device-independent cryptography had to admit that the result, as neat as it was, had no direct bearing on the device-independent scenario described above. Since Eve can manufacture the devices as she sees fit, she can make successive outputs dependent on what happened in all the previous runs.

A recent theoretical breakthrough in the analysis of the fully device-independent case, in which Eve is allowed to act arbitrarily across rounds, including through quantum entanglement – sometimes called *coherent* attacks –, is the so-called entropy accumulation theorem [ARV19; DFR20]. Roughly speaking, it allows us to "rescue" most of the above i.i.d. intuition over to the fully quantum case, giving the same bound

$$H_{\min}^{\epsilon_s}(\boldsymbol{A}|\boldsymbol{XY}E) \geq n\,\eta(S) - c\sqrt{n} \tag{7.12}$$

up to a (more complicated) correction factor c. In the later demonstration, we incorporate recent improvements [Liu+22] to the entropy accumulation theorem which increase aspects of its statistical performance.

7.3.2 *Experimental feasibility*

From the proof sketch in the previous section, the completeness condition was notably absent. It would be easy to prove by just inserting one particular model for a sufficiently well-behaved quantum device. At the same time, such a feasibility proof with a simple yes/no answer would not be very interesting in practice, as the question relevant to the practitioner is a different one: Is there a *practical device* for which a positive key length can be obtained?

To investigate this, we model the device non-idealities in an abstract, symmetric fashion by two parameters, the observed CHSH score S and the quantum bit error rate, short

QBER, Q (that is, the probability for a "bit error" causing the output of a key generation round to be anti-correlated instead of correlated). Explicitly, we define $P(A_i, B_i | X_i, Y_i)$ via

$$P(a, b|0, 0) := P(a, b|0, 1) := P(a, b|1, 0) := P(a, 1 - b|1, 1) := \frac{1 + (-1)^{a \oplus b} S/4}{4},$$

$$P(a, b|0, 2) := \frac{1 - a \oplus b + (-1)^{1 - a \oplus b} Q}{2}. \tag{7.13}$$

For this symmetric distribution of correlations in the key generation rounds, the asymptotic cost of error correction per round is simply $h(Q)$.

Looking at the difference of entropy terms in eq. (7.10), it is clear that to obtain a key of positive length, we need an experimental platform with a large Bell violation S and low quantum bit error rate Q. In practice, these will probably be somewhat correlated; for instance, if the devices are described by optimal measurements on a maximally entangled state affected by depolarising noise, then $Q = (1 - \frac{S}{2\sqrt{2}})/2$. We nevertheless keep the parameters separate, as e.g. S will in many cases be more strongly affected by decoherence than Q.

An overview of the feasibility of obtaining a positive key length in the finite-size regime for various parameters is shown in fig. 7.8 (left). Most previous detection-loophole-free Bell test experiments between two remote nodes lie in the region where it is outright impossible to obtain a key. In particular, the limited efficiency of single-photon detectors and entangled pair sources is a big challenge for purely photonic platforms, as one cannot correct for a limited detector efficiency in the device-independent context and efficient heralding schemes are not typically possible for them [Mur+19; Nie+21].

In the regime where the asymptotic key length (for $n \to \infty$) is in principle positive, finite-size statistics are another significant challenge: with the security proof being statistical in nature, an infeasibly large number of rounds might be required to obtain any key that can be certified as secret at all. This is illustrated in fig. 7.8 (right), which shows that for device parameters that are technically feasible at present, at least on the order of a million rounds are required to cross the "cliff" where the key rate l/n drops precipitously and reaches zero.

Thus, the system needs to be capable of distributing entanglement not only of high quality, but also with a rate high enough to keep the data acquisition durations in the realm of feasibility. The entanglement rate was a major challenge for past experiments using matter qubits; for instance, the S and Q link parameters achieved in ref. [Pir+10] might lie in the positive key rate region, but accumulating the $n \sim 10^8$ rounds required to obtain a key would take longer than a thousand years at the reported rate. Despite a significant amount of theoretical progress to tighten the security analysis [BHK05; Ací+07; VV14; Arn+18; Ho+20; Sch+21b; Sek+21], a practical realisation has thus far remained out of reach.
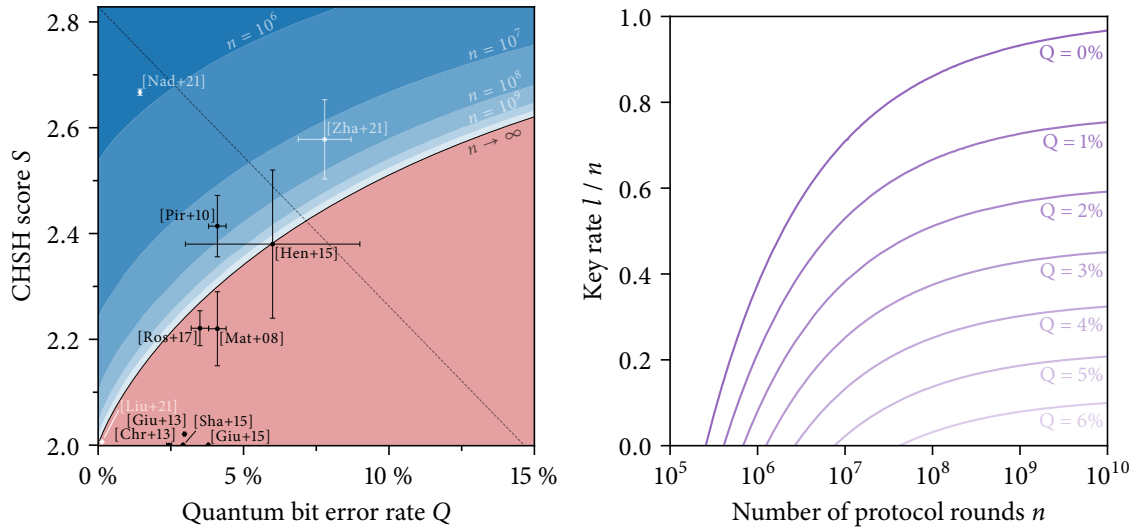
Figure 7.8: Finite-size behaviour of our DIQKD protocol for various quantum link performance models ($\epsilon_{\mathrm{snd}} = 10^{-10}$ throughout). *Left:* Requirements (in terms of S and Q) to obtain a positive-length key from a given number of rounds n. In the bottom-right red region, obtaining a key is impossible even in the asymptotic limit. As this ultimate bound is approached, the number of rounds required increases drastically due to finite-size effects; the shaded bands illustrate the decade boundaries up to which $n = 10^6, \ldots, 10^{10}$ rounds are sufficient. This is an adaption of a similar figure from a recent review for our improved finite-statistics analysis; the markers show a number of previous detection-loophole-free Bell tests compiled by its authors [Mur+19]. For many experiments the achieved data rates would be infeasibly low in the DIQKD context, as intimated by the large standard errors on the inferred S/Q parameters. Highlighted in white are three recent results, including the one presented here (top left). *Right:* Quantitative behaviour of the key length l for a number of depolarising noise scenarios (where $S = 2\sqrt{2}\,(2Q - 1)$, corresponding to the dashed diagonal in the left plot). To generate a key, $n \approx 3 \times 10^5$ rounds are required even in the noiseless case, which quickly increases for higher error rates. (Numerical calculations provided by Jean-Daniel Bancal. To fix m and ω_{thr} for a given n, a target completeness error $\epsilon_{\mathrm{com}} < 10^{-2}$ was assumed throughout, the length l of the resulting key then optimised numerically by varying γ and a number of parameters internal to the detailed proof [Nad+22].)

As discussed in chapter 6, our two-node ion-trap quantum network provides both a high fidelity and high rate of entanglement generation. As to be expected, the high Bell-state fidelity has pleasant consequences also for the S and Q performance metrics, which we shall see in the next section.

7.4   *Experimental implementation*

We implement the protocol just discussed on our two-node ion-trap quantum network. In this section, we will discuss some implementation details and performance metrics for both the data acquisition step on the quantum link, as well as the state-of-the-art classical post-processing pipeline, both of which are instrumental in successfully obtaining a positive-length key.
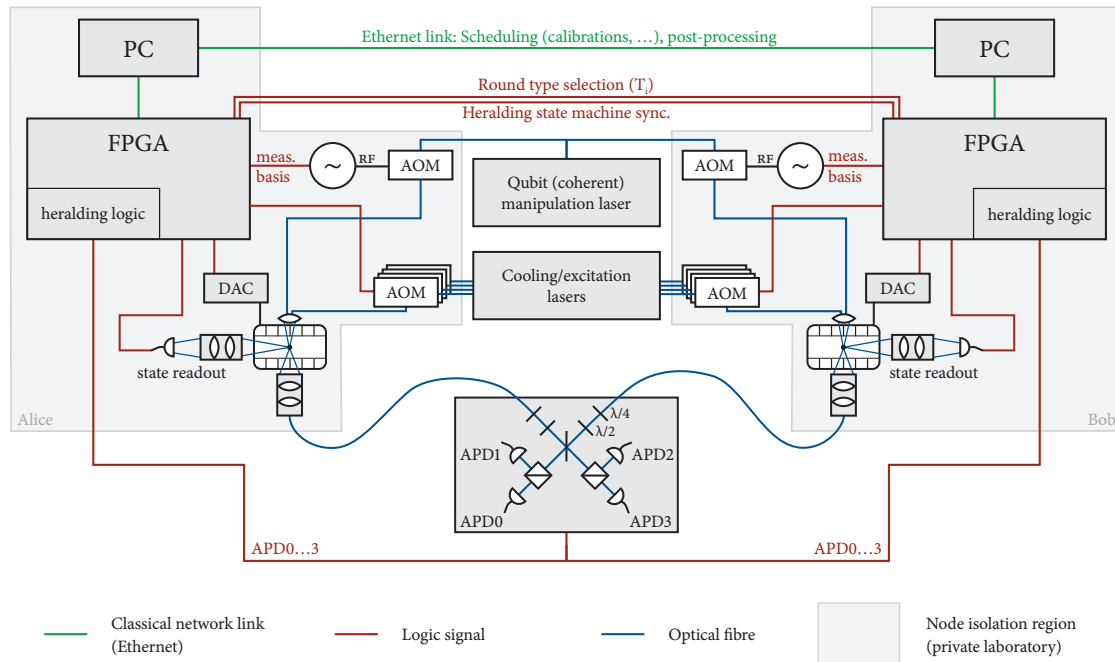
Figure 7.9: High-level sketch of the experimental apparatus viewed through the lens of a DIQKD experiment. The regions corresponding to Alice's and Bob's secure nodes are indicated by the shaded areas, their connections, including to the central heralding station, are shown explicitly. Shown are the acousto-optical modulators (AOMs) for node-local switching of the laser beams which are derived from a common set of sources, in particular also for the 674 nm qubit manipulation laser, where agile radio-frequency (RF) sources are programmed on the fly from the FPGA-based real-time control systems to choose the measurement bases. The digital-to-analogue converters (DACs) set the trap electrode voltages and are used to disconnect the photonic link by changing the ion position.

As our aim is to prove the feasibility of distributing a key with full, device-independent security using real-world devices, we opt to demonstrate all the steps as shown in fig. 7.6, and in a realistic setting: the classical experiment control systems and post-processing steps run on physically separate nodes which only communicate with each other as indicated by the protocol, and the private data never leaves the nodes – in particular, we do not manually "peek" at data at any point during data acquisition to choose suitable protocol parameters, or the like.

To align an arbitrary choice made regarding the assignment of roles in the DIQKD protocol with another arbitrary choice regarding the heralding logic in the real-time control system, the system names "Alice" and "Bob" have been swapped throughout this chapter compared to our internal naming in the laboratory and the rest of this thesis. This is a completely incidental choice; both the underlying problem as well as the hardware are in principle symmetric.
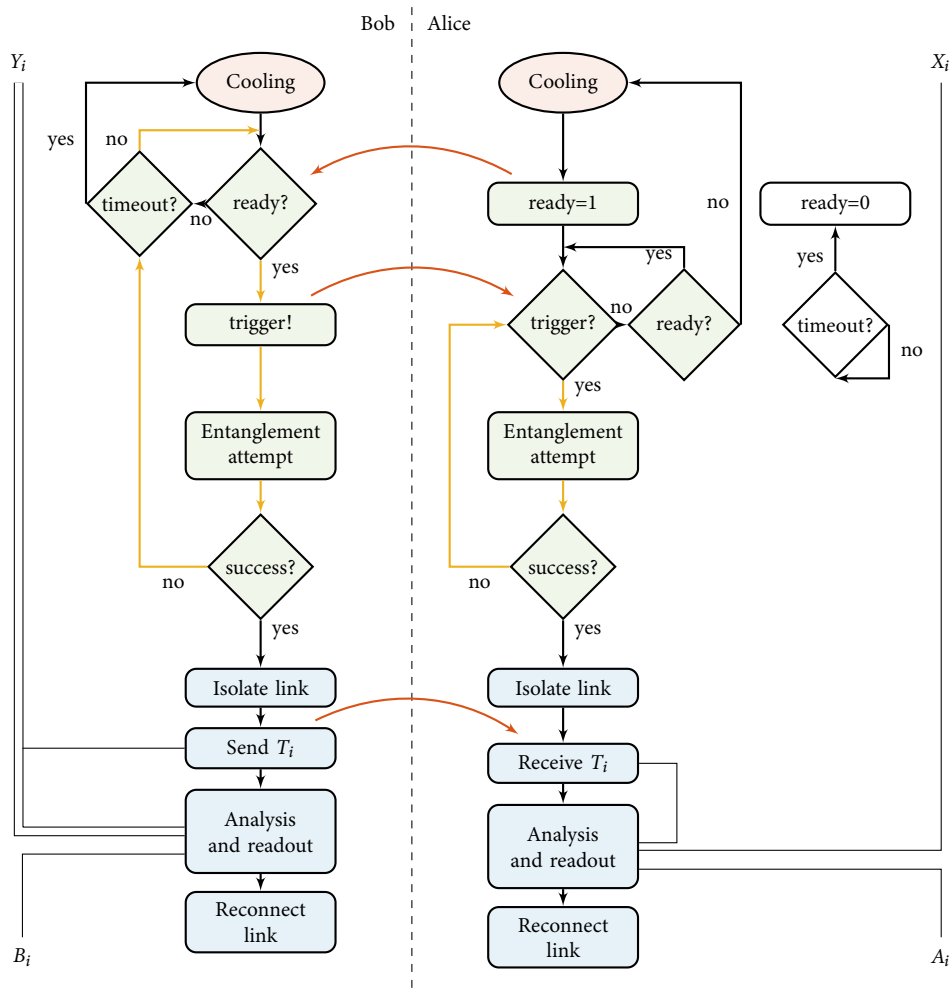
Figure 7.10: Flow chart of the real-time logic during the DIQKD data acquisition phase of the experimental sequence (left Bob, right Alice; separated by the dashed line). Execution starts on both systems with private random bits X_i and Y_i. After a period of cooling, the state machines synchronise using an acknowledged trigger pulse. Subsequently, the attempt cycle is repeated until either a successful herald is received from the external heralding station, or a set attempt duration limit has elapsed after which the ions are re-cooled. When a successful herald is received, both systems disconnect from the link by shuttling the ion away from the focus of the high-NA lens. Only then, Bob announces the round type $T_i = \chi(Y_i \neq 2)$ to Alice and executes the corresponding local analysis. Alice receives T_i and follows suit. Both parties store the outcome of their measurement, B_i and A_i, locally. After scrambling their qubit state, both parties reconnect the link, ready to start the process again for the next round. (Adapted from a diagram created by P. Drmota for the journal manuscript [Nad+22].)

7.4.1   Implementing the DIQKD boxes

The two ion-trap nodes in our laboratory are already separated into two different parts, as discussed in chapter 3, with the exception of the shared set of laser sources, which is illustrated in fig. 7.9. The latter does not compromise the device-independent security, as we will discuss in §7.6.4, so we can directly use the apparatus for the DIQKD demonstration.
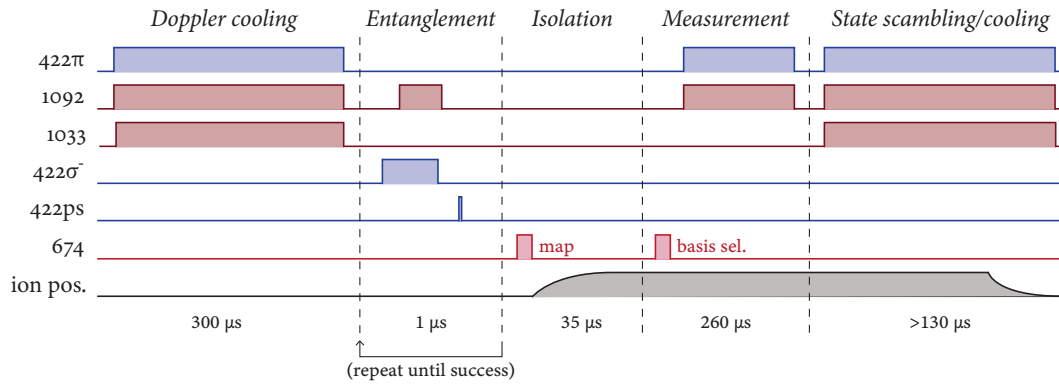
Figure 7.11: Detailed timing of the experimental pulse sequence (not to scale). The ions are Doppler-cooled prior to the entanglement attempt loop, during which tightly timed state preparation and picosecond excitation pulses create single photons entangled with the ion state. If two coincident photons are detected at the heralding station, which indicates the successful creation of a remote entangled state, the loop terminates and the nodes proceed to the rest of the sequence. The first 674 nm laser pulse is always a $\left|S_{1/2}, m_J = 1/2\right\rangle \rightarrow \left|D_{5/2}, m_J = -3/2\right\rangle$ π pulse to map the state into the optical qubit; only the second depends on the detector click pattern observed at the herald station and the selected measurement basis. The bottom trace shows the ion position along the trap axis, which is shifted to disconnect the link during the entire measurement process.

The data acquisition portion of the protocol proceeds in much the same way as the remote entanglement tomography experiment discussed in §6.3, only that the basis choice for each measurement is now dictated by a stream of node-local private randomness, and that the nodes additionally shuttle the ions away from the focus of the high-NA imaging system to prevent leakage of information outside the node during measurement. For the sake of concreteness, the heralding and data acquisition logic for the DIQKD case is shown in fig. 7.10, which includes the communication of the round type bit T_i, which Bob transmits to Alice in real-time via a shared TTL data line connecting the FPGAs.

The detailed timing of the pulse sequence is shown in fig. 7.11. After a herald event, each node first maps the entanglement from the Zeeman-split $S_{1/2}$ level onto the less magnetically-sensitive 674 nm qubit $\left|S_{1/2}, m_J = -1/2\right\rangle \leftrightarrow \left|D_{5/2}, m_J = -3/2\right\rangle$ by means of a π pulse on the $\left|S_{1/2}, m_J = 1/2\right\rangle \rightarrow \left|D_{5/2}, m_J = -3/2\right\rangle$ transition. After this, the nodes apply a step to the DC electrode voltages to shift the ions $\approx 3\,\mu\text{m}$ along the trap axis in order to disconnect the optical link, which is discussed in more detail in the context of the isolation assumption in §7.6.2.

After a delay of $\sim 30\,\mu\text{s}$, which allows the electrode voltages (and hence ion position) to settle, Alice and Bob choose the measurement basis according to their respective classical inputs X_i or Y_i via a second 674 nm laser pulse. The pulses corresponding to the settings $X_i \in \{0, 1\}$ for Alice and $Y_i \in \{0, 1, 2\}$ for Bob are chosen such that the results $A_i, B_i \in \{0, 1\}$ from the subsequent measurements fulfil the protocol requirements in the honest case, where the quantum link is not perturbed by an adversary.

In particular, for the case $(X_i, Y_i) = (0, 2)$, the key generation rounds, the outputs

should be maximally correlated, that is, the quantum bit error rate $Q = P(A_i \neq B_i | X_i = 0, Y_i = 2)$ should be minimal. As shown in chapter 4, the post-herald state for Alice's and Bob's qubits is close to

$$|\psi_k\rangle = \frac{1}{\sqrt{2}} \left(|01\rangle + e^{i\varphi_k} |10\rangle \right), \tag{7.14}$$

where the phases φ_k between the components depend on the heralding detector click patterns $k \in \{0, 1, 2, 3\}$. If we really had this, maximally entangled, state, there would be an infinite number of measurement basis pairs for which the outcomes are perfectly correlated. In the actual implementation, however, we have reason to believe that the correlations should be maximal in the computational basis. As long as the heralding station wave-plates are set such that the polarising beamsplitter axes are aligned with the π and σ decay channels, a coincident detection of two photons of opposite polarisation will always leave the ions in some sort of state in span$\{|01\rangle, |10\rangle\}$ (in the absence of any polarisation mixing effects), regardless of any dephasing that might affect the ion qubits. Furthermore, measurement in a different basis would require further single-qubit gates, which will introduce a further (small, but non-zero) error.

As the ion state is natively anti-correlated, we choose to just classically invert Bob's outcomes throughout, such that a bright measurement corresponds to $B_i = 1$, a dark measurement to $B_i = 0$. We can thus just implement $X_i = 0$ and $Y_i = 2$ as measurements in the computational basis without further manipulation of the qubit; in practice, we apply an additional 674 nm π pulse to shelve any population that might have been left over in $|S_{1/2}, m_J = 1/2\rangle$ to the additional dark state $|D_{5/2}, m_J = 3/2\rangle$ (effectively squaring the error from any mapping pulse imperfections).

What remains is to choose settings for $X_i = 1$, $Y_i = 0$ and $Y_i = 1$ that maximise the CHSH score, that is, the winning probability $P(A_i \oplus B_i = X_i \cdot Y_i \mid X_i, Y_i \in \{0, 1\})$. As our states are almost maximally entangled, we simply choose the ideal measurements that were discussed to saturate Tsirelson's bound in §7.1.1: For $X_i = 1$, Alice applies a $\pi/2$ pulse along the x axis, $R_{xy}(\pi/2, 0)$, and hence measures along the y axis. Bob instead applies one of two $\pi/4$ pulses with opposite phase, $R_{xy}(\mp\pi/4, \alpha_k)$. Evaluating this choice of projectors on the ideal state (7.14), we obtain

$$S = \sqrt{2}(1 + \cos(\alpha_k - \varphi_k)), \tag{7.15}$$

which assumes the maximum $S = 2\sqrt{2}$ if the measurement pulse phase α_k is chosen to match the state phase φ_k for the respective heralding pattern. For the case $\varphi_k = 0$, this corresponds to Alice's and Bob's measuring in the same plane on their Bloch spheres.

The experimentally observed CHSH score for a number of choices for Bob's phase parameter α is shown in fig. 7.12, showing the sinusoidal dependence expected for a state of the form (7.14). This signal could be used to empirically optimise the basis choice, although in practice, we just use the optimum predicted from high-resolution ion–ion state tomography data recorded during automated nightly calibration runs.
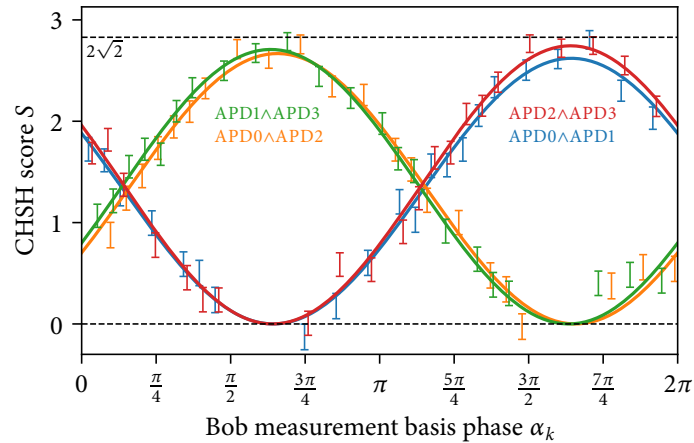
Figure 7.12: Observed Bell inequality violation as a function of measurement basis phase offset between Alice and Bob, for each of the four different two-photon detector coincidence patterns that herald the creation of remote entanglement. For each of the 21 choices of phase offset α, 5000 shots were taken, with heralds divided roughly equally between the four possible patterns; the error bars give the 68 % Clopper–Pearson confidence intervals for the binomial distribution of won CHSH games at each point. The lines show a sinusoidal fit with amplitude and phase offset floated, to guide the eye. (The horizontal misalignment in sampling grids between the four patterns is because a compensating phase offset had already been applied during data acquisition.)

A digression: "optimised" measurement angles. One possible way of cross-checking the heuristic conclusion of measuring in the computational basis being optimal would be to inspect the density matrices produced from ion–ion state tomography with sufficient resolution. Applied to the tomography data presented in §6.3, this approach however produced a somewhat unexpected result.

Performing a numerical optimisation of both Alice's and Bob's local measurement axis on the Bloch sphere to maximise correlations for each of the four heralding detector click patterns predicts quite substantial improvements in QBER, from 2.1(2) %, 1.8(2) %, 1.5(2) % and 1.6(2) % along the z axes (in the computational basis), to 1.9(3) %, 1.0(3) %, 0.8(3) % and 1.1(3) %, respectively. As illustrated in fig. 7.13, these optima occur for quite substantially tilted measurement axes, around 20° with regard to the z axis. If this could be confirmed in the experiment, the optimisation procedure could be extended to include the other measurement settings to maximise the CHSH score as well.

However, we were not able to observe this improvement in correlations when applying the requisite extra pulses to implement the tilted measurement bases in the experiment. The reason for this is unclear. As can be seen from the bootstrapped confidence intervals on the above QBERs, this did not appear to be purely a statistical artefact from the tomography reconstruction. The extra 674 nm laser pulses to rotate the state into the desired basis will introduce some extra error, but for a typical randomised benchmarking error per Clifford group element of $< 1 \times 10^{-3}$, this should be small, and should furthermore affect the tomography experiment as well. Furthermore, it is not clear where a choice of preferred angle in the $x - y$-plane would emerge from. The tomography data used here was acquired
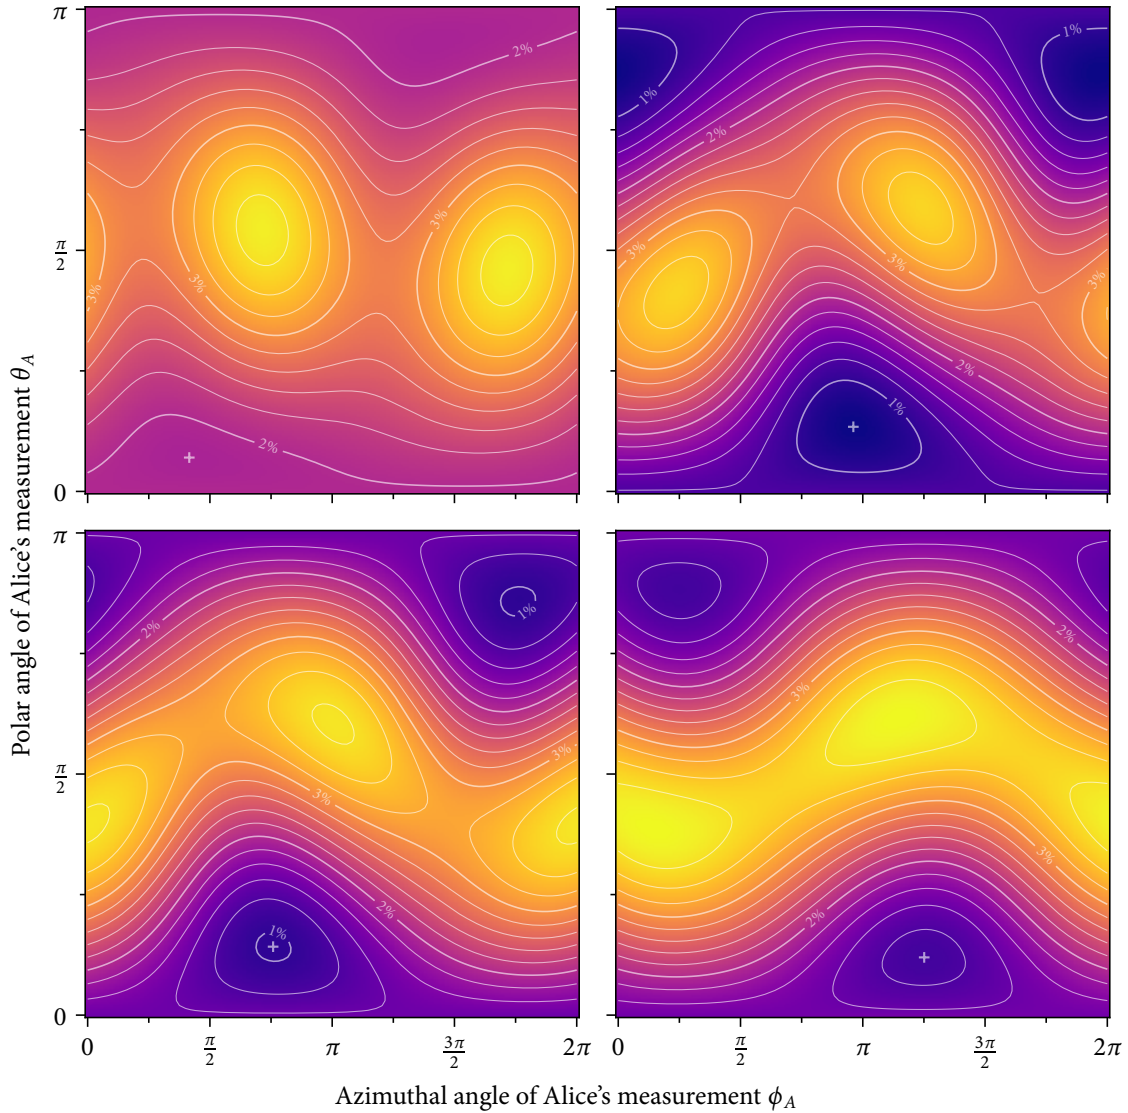
Figure 7.13: Optimal correlations between Alice's and Bob's outcomes as predicted from ion–ion state tomography data (§6.3) for each of the four detector click patterns. Each point in the plot corresponds to one choice of Alice's measurement axis; she projects onto $\left|\psi(\theta_A, \phi_A)\right\rangle = \cos\frac{\theta_A}{2}\left|0\right\rangle + \sin\frac{\theta_A}{2}e^{i\phi_A}\left|1\right\rangle$ and its complement. The colour and contour lines indicate the best QBER achievable when then optimising over Bob's measurement angles; the colour scale is the same across the plots and ranges from $Q = 0.81\,\%$ to $Q = 3.77\,\%$. Crosses mark the optima, where QBERs of 1.9(3) %, 1.0(3) %, 0.8(3) % and 1.1(3) % respectively, are expected for the four patterns (left to right, top to bottom: APD0 ∧ APD1, APD1 ∧ APD2, APD1 ∧ APD3, APD2 ∧ APD3).

X_i \ Y_i	0	1	2
0	$\begin{pmatrix} 0.415(1) & 0.0688(5) \\ 0.0961(6) & 0.420(1) \end{pmatrix}$	$\begin{pmatrix} 0.439(1) & 0.0805(6) \\ 0.0735(5) & 0.4068(10) \end{pmatrix}$	$\begin{pmatrix} 0.5017(5) & 0.00339(6) \\ 0.0110(1) & 0.4839(5) \end{pmatrix}$
1	$\begin{pmatrix} 0.3928(10) & 0.0916(6) \\ 0.0851(6) & 0.431(1) \end{pmatrix}$	$\begin{pmatrix} 0.0820(6) & 0.437(1) \\ 0.3970(10) & 0.0841(6) \end{pmatrix}$	—

Table 7.1: Empirical probabilities to observe classical outcome (A_i, B_i) for measurements settings (X_i, Y_i). Matrix rows index Alice's outcomes, columns that of Bob. The probabilities are estimated from the 1 920 000 total characterisation rounds also shown in fig. 7.15 and fig. 7.14, with the multinomial standard errors given in parentheses. Bob's classical outcomes were inverted (such that 1 corresponds to finding the ion in the $S_{1/2}$ manifold after shelving), giving maximum correlations for the key generation settings $X = 0, Y = 2$, and maximising the probability that $A_i \oplus B_i = X_i \cdot Y_i$ in Bell test rounds.

with an over-complete set of measurement settings corresponding to $\{Z, X, Y, -X, -Y\}$, which should prevent gate and readout errors from introducing such a bias.

At the time of writing, this discrepancy has not been resolved yet; one possible mechanism would be a low-level drift in birefringence in the single-photon paths, which would require a more systematic investigation to track down. For DIQKD, the QBER when just measuring in the computational basis turned out to be sufficient for our demonstration. However, care should be taken when using the tomographic data for other estimates sensitive to the exact level correlations along a particular axis, such as possibly for entanglement distillation.

7.4.2   Link performance

With this implementation of the two DIQKD "black box" primitives (classical input → classical output) in hand, we first evaluate the performance of the remote entanglement link, both to later be able to choose appropriate protocol parameters for the DIQKD demonstration experiment, but also to document its performance, as none of that information is revealed by the actual protocol (beyond the pass/fail information on Bell test verification and error correction). For this purpose, we run the data acquisition portion of the protocol exactly as we would for the demonstration experiment, but with approximately equal amounts of Bell test and key generation rounds ($\gamma \approx \frac{1}{2}$) to gather statistics on both round types.

The most expansive data run was taken on the same weekend as the final DIQKD demonstration, and comprises $n = 1\,920\,000$ total rounds. After completion, the data from each node is centrally combined and interactively analysed (in contrast to the actual DIQKD run!). The raw results are shown in fig. 7.14. As expected for a reasonably well-behaved experiment (without, for example, undetected periods of complete equipment failures), the outcomes visually look fairly uniformly distributed, without any obvious aggregations of errors or the like.

| | Key generation round, both 0 | | Key generation round, mismatch | | CHSH game win, Alice 0 |
| | Key generation round, both 1 | | CHSH game loss | | CHSH game win, Alice 1 |
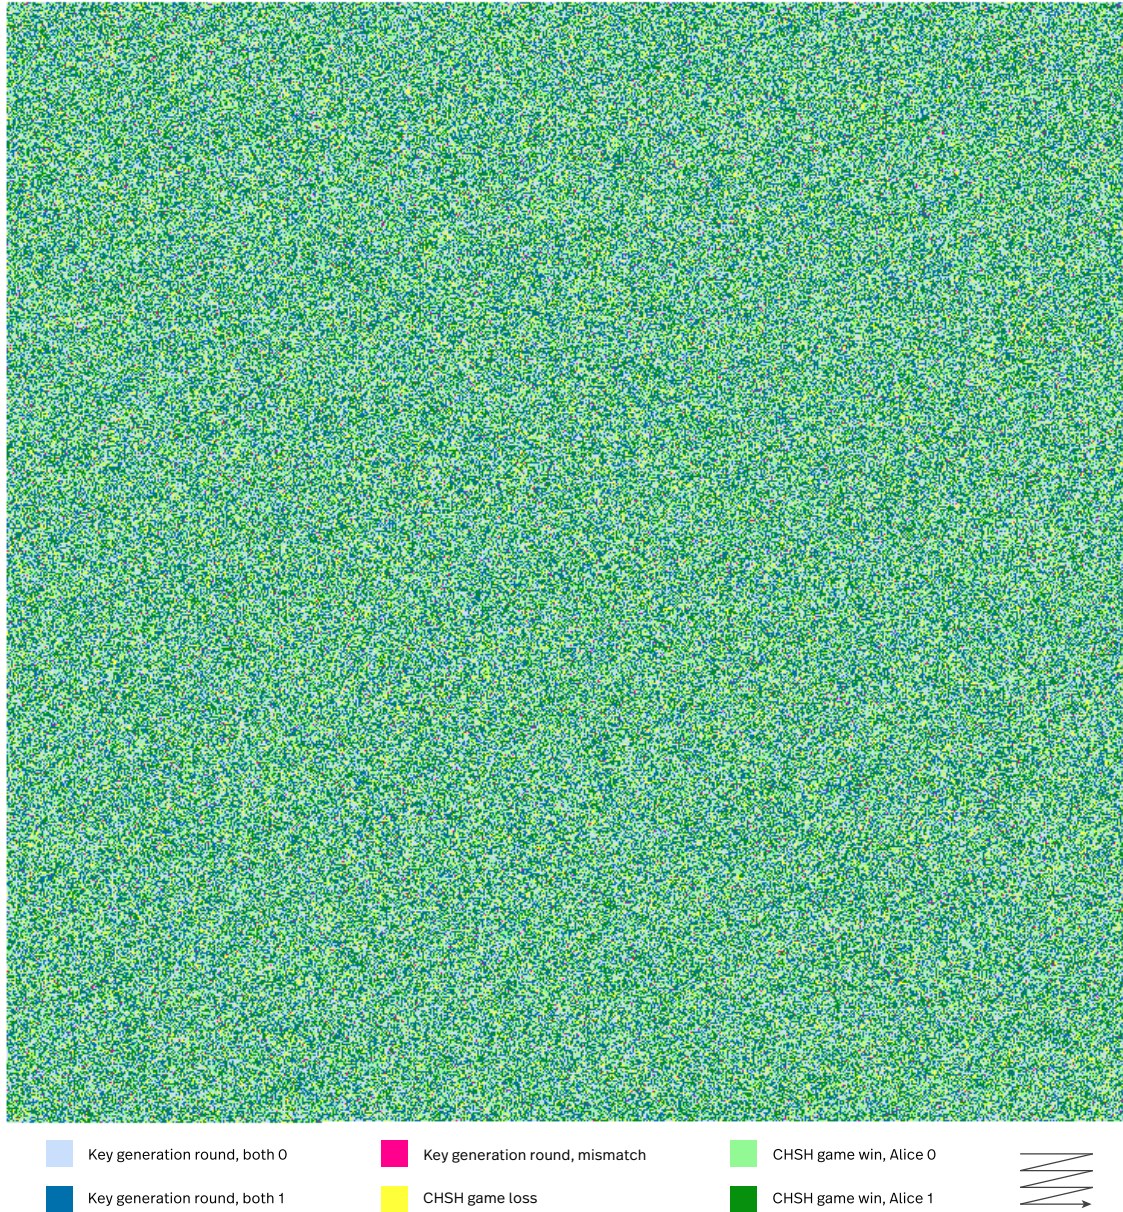
Figure 7.14: Round-by-round outcomes for the DIQKD link characterisation run. Each square shows the round type ($\gamma \approx 1/2$) and outcome for one of 1 920 000 rounds in the dataset, arranged row by row from left to right and top to bottom as indicated by the arrow. The experimental sequence is the same as for the actual DIQKD demonstration, but only here, the classical outputs are collected from the node PCs afterwards to create the visualisation.
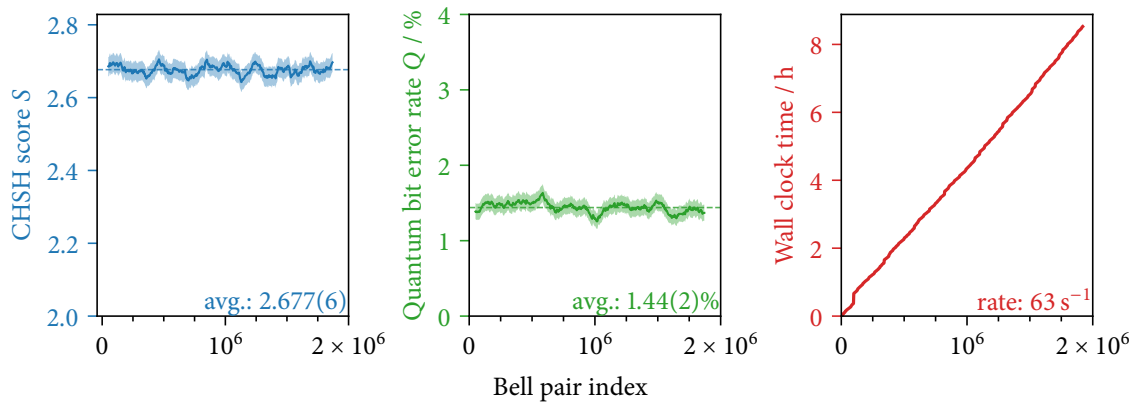
Figure 7.15: A separate characterisation run of 1 920 000 total Bell pairs shows stable link performance. Inputs and outputs from both nodes were collected to compute a moving average for the CHSH score S and the quantum bit error rate Q (window length: 100 000 rounds; test round fraction: $\gamma = 1/2$). The shaded bands indicate 95 % confidence intervals from binomial statistics in each window. The rightmost panel shows the acquisition timestamp for each Bell pair during the 8.5 h experiment duration. The vertical steps, where time passed without heralds, correspond to ion loss and recalibration events..
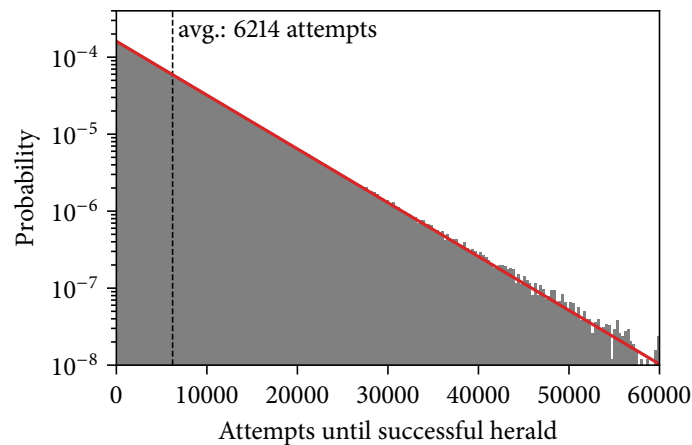
Figure 7.16: Histogram of the number of excitation attempts necessary until a two-photon herald event was observed during acquisition of the link performance characterisation data. Data from 1 920 000 total heralds is sorted into 200 bins; the solid line shows the exponential distribution corresponding to a Poisson process with success probability $1/6214 = 1.61 \times 10^{-4}$. The effective attempt rate after including the amortised cost of re-cooling for 250 μs every 1 ms is 800 kHz, yielding a raw entanglement rate of 129 s^{-1}.

A more easily interpreted summary of the data is shown in fig. 7.15, where the CHSH score and quantum bit error rate were computed in a moving window, and can be seen to be approximately constant across the entire experiment duration. The full description of the "black boxes" in terms of the empirical probabilities to observe certain outputs for a given pair of inputs, averaged over all the data, is shown in table 7.1.

Figure 7.15 also shows the time at which each Bell pair was generated, relative to the start of the data acquisition process. Overall, the data acquisition rate is fairly constant, and the "raw" data acquisition rate of 129 s^{-1} (see fig. 7.16) actually exceeds the 100 s^{-1} reported in chapter 6 (which was accidentally limited by a short attempt duration between

$$P_{T_i=1}(A_i, B_i) = \begin{pmatrix} 0.4210 & 0.0807 \\ 0.0847 & 0.4136 \end{pmatrix} \qquad P_{T_i=0}(A_i, B_i) = \begin{pmatrix} 0.5017 & 0.0034 \\ 0.0110 & 0.4839 \end{pmatrix}$$

Table 7.2: Channel statistics assumed for the Bell test rounds (left) and key generating rounds (right) for error correction; these probabilities define the decoder priors. The matrices give the expected correlations between each bit in Alice's and Bob's strings $A = (A_i)_i$, $B = (B_i)_i$ (after flipping Bob's result for the $X_i = Y_i = 1$ Bell test case), with rows corresponding to Alice's outcomes and columns to Bob's.

interruptions for recooling). The average rate across the entire duration of the experiment is significantly lower, however, at $63\,\text{s}^{-1}$. The biggest limitation was that the ion lifetime in Alice had fallen to $\sim 5\,\text{min}$ due to a vacuum leak, which limited the achievable duty cycle due to the need to frequently re-load the trap, and thus the overall rate across long times[7]. These periods of the nodes waiting for Alice's trap to reload an ion, along with other periodic recalibration events (e.g. of the single photon fibre alignment), are visible in fig. 7.15 as vertical steps in the acquisition time curve.

7.4.3 Classical post-processing

The implementation details of the classical post-processing pipeline matter greatly for the feasibility of the experiment, as the amount of information revealed by an uneconomic implementation during reconciliation could eclipse the amount of private entropy certified from the given Bell inequality violation, possibly leading to an overall negative key rate. This point, the "frugality" of the used algorithms, thus received significant attention by our theory collaborators. Here, we only discuss a few salient points, in particular as they relate to the link performance, and consequently, the parameter choices made for the experiment.

Error correction. One challenge in a practical implementation of DIQKD is in the error correction process, as it is imperative to keep the overhead as low as possible, as it directly reduces the key rate. To that end, many previous theoretical analyses [ARV19; Mur+19; Tan+20] assumed codes that can only be decoded by effectively inverting a hash function, which is computationally infeasible. In our work, we use a spatially coupled low-density parity-check (SC-LDPC) code to achieve performance close to the theoretical limit – LDPC codes are known to be (Shannon-)capacity-approaching, and SC-LDPC codes admit efficient decoders that converge to the maximum-likelihood decoding threshold of the underlying LDPC code. The code was designed and implemented by K. Ivanov, and is described in more detail in his doctoral thesis [Iva22, ch. 5].

For decoding, we implement a belief propagation decoder based on the standard sum-product algorithm. This is a "soft-decision decoder" that iteratively approximates the

---

[7]This severe duty cycle limitation was the reason why we opted not to optimise the system efficiency at the time, even though we could have presumably reproduced or exceeded the $182\,\text{s}^{-1}$ rate we reported in ref. [Ste+20] without much difficulty.

maximum-likelihood decoding through a series of updates for the likelihood distribution of each bit. This allows us to trivially incorporate the information about the different statistics in Bell test and key generation rounds into the prior distributions assumed for each bit, and treat the entire output strings in one go, which turns out to be advantageous for the overall entropy cost of reconciliation. For this, we assume the distributions shown in table 7.2, obtained by averaging across the various Bell test inputs from table 7.1. The decoder is implemented in C++, and includes a simple stopping criterion, where the iterative scheme is terminated if the (hard-decoded) guess at the output does not change between successive rounds. For the $n = 1.5 \times 10^6$ rounds in the actual DIQKD demonstration, this criterion was reached in $\approx 4\,\text{min}$.

The construction used includes the possibility to "prune" some of the parity checks to fine-tune the length of the syndrome, m. To pick a syndrome length that is high enough to almost guarantee successful decoding, yet not unnecessarily wasteful in the entropy budget, we model the correlations in the experiment according to eq. (7.13) and perform numerical experiments for various numbers of rounds n. Based on this, we decide to set

$$m = \left(\gamma\, h(\omega) + (1 - \gamma)\, h(Q) \right) n + 50 \sqrt{n} \qquad (\text{where } \omega = (S + 4)/8) \qquad (7.16)$$

in the later experiments, where the extra overhead on top of the asymptotic entropy cost, expressed by the numerically determined pre-factor to the $\sqrt{n}$ term, ensures successful decoding in $> 99\,\%$ of simulations.

Privacy amplification using a strong extractor. In the privacy amplification step, we use a randomness extractor to obtain a short key from the (error-corrected) long strings of raw measurement outcomes. An extractor is a function that takes a "weakly" random string X together with a uniformly random, short seed S, producing a shorter, but uniformly random output. Crucially, if the seed S is not known to Eve, the output appears entirely uncorrelated to any partial knowledge she might have about Alice's and Bob's outcomes that are used as the input, thus producing an (almost) perfectly secure key.

We choose an extractor based on Trevisan's construction [Tre01] with a block weak design, which has been shown to be quantum-proof [De+12], and furthermore is a "strong" extractor: without complete knowledge of the inputs, the seed S appears uncorrelated to the output, enabling it to be re-used. The information-theoretical treatment is discussed in detail in supplement S9B to ref. [Nad+22].

The extractor was implemented by N. Sangouard, based on an open-source implementation by W. Maurer [Mau14]. The performance is sufficient such as not to limit the post-processing pipeline; to extract 95 884 bits from 1.5×10^6 raw input required 21 s on the node PCs. For this, an initial key of 1 201 886 bits is necessary, but is entirely reusable between rounds due to Trevisan's construction being a strong extractor (see table 7.4).

Message authentication through universal hashing. To generate short signatures of the classical communication in the protocol, as well as to verify the correctness of the error correction output (which are both bit strings with length $\sim n$, the number of rounds), an obvious choice is the use of hash functions: mathematical "one-way" functions that produce a short output from a potentially much longer input, in a fashion where the function is hard to invert. Simple hash functions such as *SHA* family of standards published by NIST are in widespread use in classical computing, but on their own are insufficient for information-theoretic security, as with the exact knowledge of the algorithm, it would in theory be possible for an attacker to find colliding inputs to spoof a message (if thought to be computationally infeasible in practice).

We thus employ universal hashing, which builds on a whole family $\mathcal{H}$ of hash functions $f_p \in \mathcal{H}$, parametrised by some seed parameter p. Colloquially, the desired universality property means that if we pick a hash function from the family at random, the possibility for two distinct inputs $X \neq Y$ to collide, that is, to produce identical outputs $f_p(X) = f_p(Y)$, is very low (ideally close to 2^{-l}). Formal definitions for this and some related properties, as required for the security proof, are recapitulated in supplement S9A to ref. [Nad+22].

Practically, this means that if Alice and Bob share a private seed p ahead of time to select a function from the family, Eve cannot generate the correct tags for any messages she were to alter. In our case, we use the *VHASH* universal hash family to generate authentication tags of 64 bit length, achieving a collision probability of $2^{-61} \approx 4 \times 10^{-19}$ [DK07]. For this, the seed p needs to be 1280 bits long. Chiefly to enable Alice and Bob to re-use this seed across the entire protocol, we then encrypt the output with a fresh one-time pad each time it is used. This procedure was implemented by J.-D. Bancal based on the reference implementation of the *VHASH* family [Kro07], and is available as *UVMAC* [Ban21]; the contribution to the post-processing time is minimal.

This cost of 64 bits of private key for each authentication tag is the only actual consumption of (non-reusable) private randomness (see table 7.3), marked in fig. 7.6 with a key symbol. As discussed in §7.2.1, some consumption of a short pre-shared key for information-theoretically-secure authentication cannot be completely avoided.

7.5   *Distribution of a key*

What is left for a full demonstration of DIQKD according to the protocol from fig. 7.6 is to choose the parameters n, γ, ω_{thr} and m, to initialise the required pre-shared key K_0, and finally to execute the steps on the two-node network. For this first demonstration, our aim was to obtain a key with full device-independent security (including a suitably low security parameter ϵ_{snd}), but without any particular focus on the key rate or other

performance metrics, as long as the length of the freshly generated key would exceed the number of pre-shared key consumed for authentication.

Previous numerical investigation of the detailed entropy calculations in the full security proof indicated that, based on some preliminary link performance characterisation data, $n \gtrsim 1 \times 10^6$ rounds should comfortably result in a non-zero key length. The choice of test round fraction γ is a trade-off between the number of rounds allocated to Bell tests and the tightness of the resulting bounds on the adversary's knowledge, and for this range of parameters was shown to result in an optimal key rate for $\gamma \approx 5\%$ (see e.g. Fig. S15 and S16 in the supplement to ref. [Nad+22]).

The final choice of parameters for the results presented in ref. [Nad+22] was made based on some preliminary link characterisation data. At this point, the large dataset from fig. 7.15 was not available yet, and the entropy calculations had not been finalised. We somewhat arbitrarily chose $n = 1.5 \times 10^6$ (expecting a manageable experiment duration of $\sim 7\,\mathrm{h}$), and a test round fraction of $\gamma = 13/256 \approx 5.1\%$ (with the rational power-of-two fraction to trivially avoid any biases in the random number generation, see §7.6.5). From the characterisation data, we decided on a conservative bound on the system performance, expecting a Bell score of $S > S_{\min} := 2.64$ or a winning probability $\omega_{\min} := (S_{\min} + 4)/8 = 0.83$. Note, however, that ω_{thr} really enters the protocol as an upper bound $n\,\gamma\,(1 - \omega_{\mathrm{thr}})$ on the allowable number of CHSH game losses observed across the finite number of test rounds. Both the number of test rounds (we expect $n \times \gamma \approx 76171$) and the number of failures among those result from probabilistic processes. In the honest implementation, each round is independent and results in a failed Bell test with probability $\lambda := \gamma\,(1 - \omega_{\min})$. To guarantee a low possibility of spurious failures due to statistical fluctuations, we further "derate" the maximum number of accepted failures by $k = 3$ standard deviations from the expected value, and accordingly choose ω_{thr} as

$$\lambda_{\mathrm{thr}} := \frac{1}{n}\left(n\lambda + k\sqrt{n\,\lambda(1 - \lambda)}\right), \quad \omega_{\mathrm{thr}} := 1 - \frac{\lambda_{\mathrm{thr}}}{\gamma} \approx 0.825\,538. \tag{7.17}$$

For choosing the length of the error correction syndrome, we similarly decided to estimate the link performance with the bound $Q < Q_{\max} := 1.8\%$, and according to the model from eq. (7.16), chose $m = 296\,517$ to ensure a low chance of spurious error correction failures ($\leq 0.1\%$) for the honest implementation. As the subsequently acquired long characterisation dataset (fig. 7.15) shows that the link performance at the time of data taking was described by a typical CHSH score $S = 2.677(6)$ and bit error rate $Q = 1.44(2)\%$, these assumption were indeed very conservative, and the protocol is expected to succeed with significantly better probability than 99% (such that failures probably are not any more likely than random failures of the laboratory equipment).

The protocol is now executed on the two separate node control computers. A set of
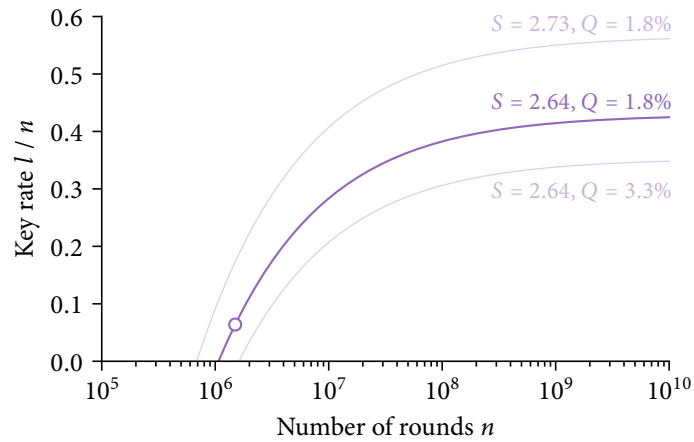
Figure 7.17: Finite-size key rates achievable for a soundness error of $\epsilon_{snd} = 10^{-10}$. The circular marker on the main curve corresponds to the parameters for the main DIQKD demonstration, where we assumed a link performance of $S_{min} = 2.64$, $Q_{max} = 1.8\,\%$, and chose $n = 1.5 \times 10^6$, $\gamma = 13/256$, $\omega_{thr} \approx 0.825\,538$ and $m = 296\,517$, resulting in a key of length $l = 95\,884$. This is a rate of $r = l/n \approx 6.4\,\%$, which reaches 42 % at $n = 1 \times 10^{10}$ rounds. The two alternate curves illustrate its sensitivity to link parameters. The choices were made to correspond to measurements on a depolarised state with the same S and Q values as S_{min} and Q_{max}, respectively, illustrating the impact of the better correlations in the computational basis observed in the experiment. (Calculations supplied by Jean-Daniel Bancal.)
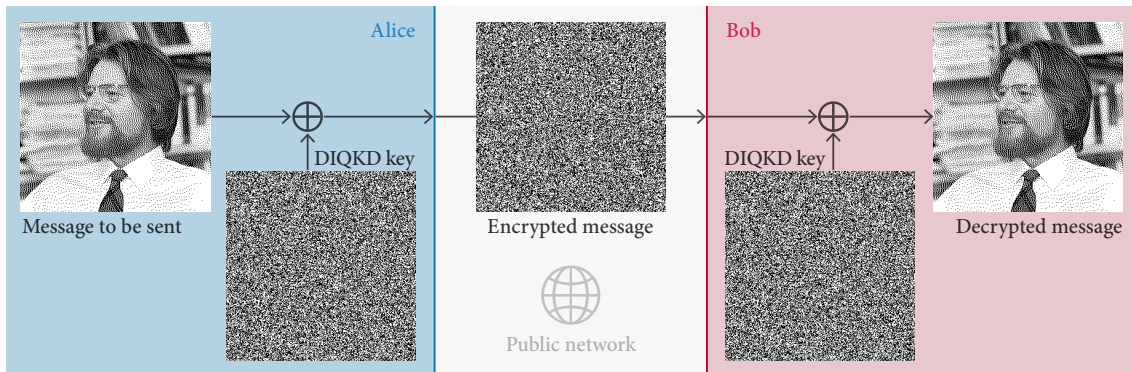
Figure 7.18: Once Alice and Bob have obtained a secret key through DIQKD, they can use it as a one-time pad for provably secure communication. To illustrate this in the experiment, Alice transmits to Bob an encrypted picture of John S. Bell (original photo: CERN). To encrypt the 309 pixel × 309 pixel black-and-white picture, 95 481 bits of the generated key are used, and an additional 64 bits to encrypt an authentication tag for the encrypted message to guarantee integrity. The protocol generated 95 884 bits, of which 256 additional bits are required to replenish the consumed private key (see table 7.3) to be able to execute the DIQKD protocol again, leaving 82 bits of additional unused key. This demonstration was executed at the very end of the project, and is the first time any of the secret information was used or inspected.

Privacy amplification	
Certified fresh entropy	95 884 bits
One-time pads	
error correction	−64 bits
auth. Bob to Alice	−64 bits
auth. Alice to Bob	−64 bits
key activation	−64 bits
Overall key length gain	95 628 bits

Universal hashing seed	1280 bits
Trevisan's extractor seed	1 201 886 bits
Total reusable secrets	1 203 166 bits

Table 7.3: Entropy balance sheet for the DIQKD demonstration experiment. Based on the chosen parameters n, γ, ω_{thr} and m, if the protocol succeeds, we can extract 95 884 bits of uniformly random, private key with $\epsilon_{\mathrm{snd}} = 10^{-10}$. In the process, the protocol consumes 256 bits from the initial key K_0 for authentication, leading to an overall increase of $|K_1| − |K_0| = 95\,628$ bits.

Table 7.4: Reusable private randomness required to be shared between Alice and Bob for the protocol. As we use a strong extractor for privacy amplification, and all the authentication tags are encrypted using a one-time pad, no information leaks to the adversary. Hence, this key, while significantly longer than the amount of entropy produced, can be reused entirely between rounds.

simple Python scripts drive the data acquisition[8], and connect the various post-processing steps, as discussed in the previous section. The outcome strings, along with other private data, neither leave the respective node PCs, nor are they manually inspected; the public messages required by the protocol according to fig. 7.6 are exchanged over the Ethernet network[9].

We acquire data over 7.9 h of wall-clock time, interrupted for 4.4 h due to an unexpected laser failure[10]. As required for security, a heralding event during protocol execution always resulted in a bit added to the final A and B outcome strings; no post-selection took place. Choosing a soundness error bound of $\epsilon_{\mathrm{snd}} = 10^{-10}$, the detailed security analysis (see supplements S10 and S11 to ref. [Nad+22], particularly eq. (73)) yielded an extractable key length of 95 884 bits, equivalent to a key rate of $l/n \approx 6.4\,\%$ as illustrated in fig. 7.17. The classical part of the protocol completed successfully over a total duration of 5 min, indicating that Alice and Bob share a random key that is, with very high probability, private and identical. As the amount of new key material greatly exceeds the 256 bits of private randomness consumed during the protocol (see table 7.3 and table 7.4), Alice and Bob can extend their private keys by 95 884 − 256 = 95 628 bits. Thus, this marks the first demonstration of fully device-independent quantum key distribution.

[8]The high-level implementation of the post-processing pipeline was developed in collaboration with Jean-Daniel Bancal.

[9]For simplicity, this was implemented using the *scp* file copy utility, which happens to be based on the encrypted *SSH* (Secure Shell) protocol. Of course, this encryption layer is irrelevant here; a simple plain-text socket connection could have been employed instead.

[10]More precisely, this was a sporadic RF DDS failure that caused no 674 nm power to be available at one of the nodes, and caused execution to halt after it was detected by the closed-loop power stabilisation system. As this is in the part of the system that is "self-tested" via the observed CHSH score, this is not a detriment for the security of the resulting key – only the experimentalist's pride, as other, similarly long experiments successfully completed without any such issues.

We made no attempts to characterise the long-term stability of the protocol with these parameters, but once again stress that they were chosen quite conservatively – had the aim been to optimise the key rate, rather than to just demonstrate the feasibility of DIQKD on real-world hardware, more aggressive parameters could have been chosen (and before all, a larger number of rounds n). On the same weekend that the characterisation data and DIQKD results presented here were obtained, two more protocol runs were taken, one with $n = 1 \times 10^6$, $S_{min} = 2.64$, $Q_{max} = 2\%$, and another with $n = 2 \times 10^6$, $S_{min} = 2.65$, $Q_{max} = 1.7\%$. The error correction verification step – the last chance for failure in a scenario where the classical communication is not tampered with – succeeded for both these runs. Although we opted to use the $n = 1.5 \times 10^6$ run to present as our main result, keys could have been extracted from those runs as well.

7.6 *Plausibility of assumptions*

While our comprehensive security analysis shows that the resulting key is secret against the most general quantum attacks (and with a security parameter small enough to be beyond controversy), this high level of security critically rests on the DIQKD assumptions described in §7.2.1. The isolation assumption (ii), requiring that no unwanted information can leak *out of* Alice's and Bob's nodes, deserves particular scrutiny, as it must hold against a technologically arbitrarily capable attacker outside Alice's and Bob's laboratory.

Some amount of dependence on the hardware details of the quantum devices is unavoidable. We do not concern ourselves here with the complete shielding of electromagnetic emanation from the two nodes. While an issue that requires care in practice (as evidenced by the various standards for shielding of environments in which classified information is handled), our DIQKD setup is not substantially different from any other telecommunication/computing use case, as the necessary laser technology can be kept locally at each node. What remains to discuss is the quantum link, where we, in practical terms, need to assume that the adversary can at will direct the fibre outputs onto a perfectly efficient detector, and can also shine arbitrarily powerful laser beams back into the system.

Some of these concerns are shared with Bell test experiments used to disprove local realism, where deviations between the underlying theory of reality, the model describing the experiments, and the apparatus itself are discussed under the term "loopholes". As DIQKD also incorporates a CHSH inequality violation, we will first briefly discuss some parallels and differences between the settings, before turning to some data which makes it plausible that our technical realisation of the isolation mechanism and random number generation should be trustworthy.

7.6.1   Bell-test loopholes: Space-like separation vs. locality

The aim of experiments colloquially referred to as "Bell tests" is to show that the behaviour of Nature is inconsistent with a realist local description, often by observing a CHSH score violating the inequality $|S| < 2$ which holds for all such theories [Bru+14]. A series of early experiments (e.g. ref. [AGR81]) provided increasingly convincing evidence that this should be the case, but suffered from a number of loopholes that would still permit a local description if different parts of the apparatus suitably "conspired" against the experimenter. Two considerations stand out among this series of loopholes, which were only finally all addressed at the same time by a series of celebrated experiments in 2015 [Giu+15; Hen+15; Sha+15]: the *detection* or *fair-sampling* loophole, and the *locality* loophole.

The detection loophole is a consequence of the finite quantum efficiency of (typically) single-photon detectors. It arises if that limited efficiency is then backed out from the observed data by assuming that it equally affects all measurements, that is, that the observed events are a fair sample of an underlying true behaviour (independent of e.g. the measurement settings). While the assumption that Nature does not actively conspire against the experimenter to modulate the detector loss in complex ways might be defensible in a Bell test experiment, clearly the detection loophole is of grave concern to DIQKD, where we explicitly assume the possible presence of a crafty, active eavesdropper. As discussed in §7.4.1, we thus avoid post-selection, that is, we commit to appending the measurement outcomes to the **A** and **B** strings as soon as we observe a herald signal – and crucially, before revealing the round type or applying the measurement settings to the apparatus.

The Bell test and DIQKD scenarios markedly differ regarding the locality loophole, however. If, in a Bell test experiment, the measurement setting choice made on one side could propagate to the other node and influence the measurement outcome there, a realist local explanation would again be possible. In order to exclude the possibility of such information transfer, even through some kind of yet-unknown physical mechanism, Bell test experiments can be performed such that the measurement on one node is not in the relativistic light-cone of the random measurement basis choice on the other node, hence guaranteeing the absence of timely communication through space-like separation. Once the classical measurement results are determined, there is no harm in information transfer between the nodes; in fact, all the input and output data from the nodes is later brought together to compute the CHSH score S (in much the same way as in §7.4.2). Space-like separation of one node from the random basis choice at the other only holds for a limited amount of time, but this is (with clever engineering) sufficient for a Bell test experiment.

The setting of DIQKD, however, is fundamentally different. While the measurement settings must also be unknown to the remote side in our case, this also applies to the results, i.e. the classical outputs – and crucially, these need to remain private for all times. If the outputs do eventually reach an adversary, the generated key material is compromised. Whether this leakage process is restricted to subluminal speeds or not is of little

help in keeping the information secure; it needs to be excluded entirely. Either way, it is not clear what would even be meant by "space-like separation" in the case of a DIQKD protocol like described in fig. 7.6, which of course includes interactive communication between the nodes. But even in a hypothetical protocol that is executed instantly and does not require any classical messages for the reconciliation process, speed-of-light considerations could only guarantee the security of a key for at most $\approx 43\,\text{ms}$ (if the adversary is known to reside on the opposite side of Earth's surface).

More generally, DIQKD is a cryptographic protocol, and as such operates in the setting of cryptography, where the ability to protect private information from leaking to the outside is a foundational prerequisite, and already required just for the term "private key" to be well-defined and meaningful. Since the isolation condition by itself already guarantees locality for the Bell-test component of the protocol, space-like separation of any parts of the protocol does not help any further.

We thus conclude that space-like separation of some (which?) protocol events is neither necessary nor sufficient for security of a DIQKD protocol, which is well-understood in the quantum information theory community [Mur+19; Pir+10; Pir+09; Pir+20]. Rather, the isolation of the nodes necessarily remains as an assumption, which has to be ensured through technical means.

7.6.2 *Shuttling-based isolation*

Because the 422 nm transition used for remote entanglement generation is the same as is used for the state-dependent fluorescence readout, a naive implementation of the DIQKD primitives in a trapped-ion remote entanglement experiment would, with high certainty, leak the outcome of the measurement out of the nodes. As mentioned above, it is not enough to consider the photons emitted during the intended readout process; the attacker could actively couple resonant laser light back into the fibre leaving the node to excite the transition outside that period.

This is not an issue during or immediately after the remote entanglement generation process, as such attacks would in this case be detected by a corresponding drop in observed CHSH score. The critical section starts with the (optional) pulse that implements the basis change for the Bell test rounds, and ends only with the reinitialisation of the qubit state (see next section). During this period, the qubits need to be well-isolated from the optical fibre link. As this is in tension with the requirement for the qubits to be well-coupled to the link during the heralded entanglement generation step, an active disconnection mechanism is needed. Because of the good coherence times achievable in trapped ions, it is not particularly challenging to implement this. In a real deployment, the isolation could for instance be established to an effectively arbitrary degree through a free-space shutter or a fast micro-electromechanical (MEMS) fibre switch, which would have the advantage of being a macroscopic component that is easy to separately verify. Another option would be
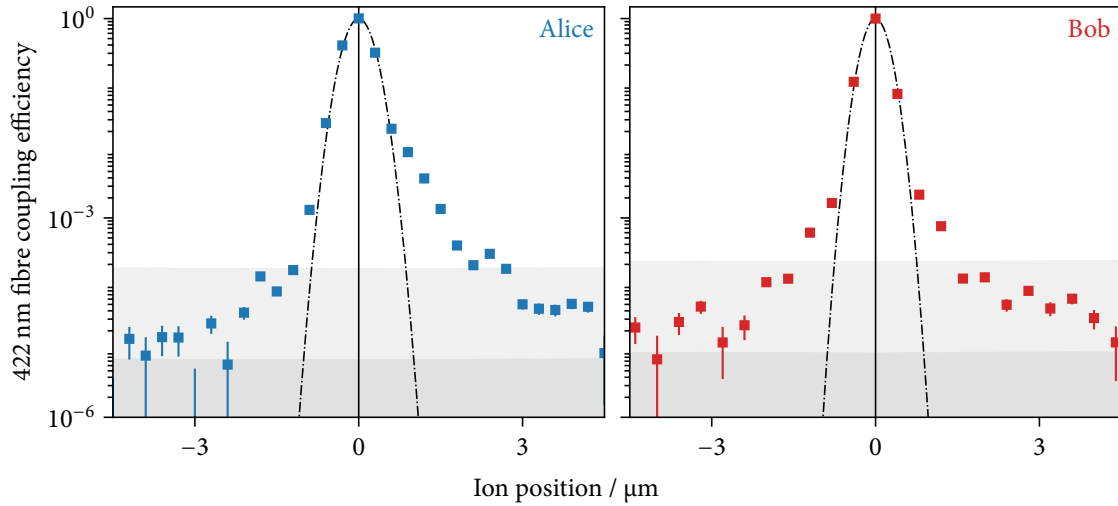
Figure 7.19: Position-dependent coupling efficiency from ion to fibre link, normalised to that in the trap centre. For each point, detector counts were integrated over 30 s with the ion at the specified position (through a trap electrode voltage set calculated from the electrostatic trap model) and the Doppler cooling beams turned on. The positional variation in fluorescence was corrected for by normalising against the respective count rates on the position-insensitive state readout PMT, although the variation in beam intensity over the region shown here is small (< 4 % and < 8 % for Alice and Bob, respectively). The parabolas show a least-squares fit of a Gaussian to the data; the actual behaviour from a convolution of imaging point spread function and fibre mode is more complex, and in addition, heavily affected by aberrations in the wings. The upper curve (approximately a horizontal line) up to which the background is shaded corresponds to the counts from the $18 \, \mathrm{s}^{-1}$ dark count rate of the used detectors, which were subtracted from the data. The lower curve up to which the background is shaded more intensely represents the uncertainty in the background count rate measurement (assuming Poissonian statistics), which limits the dynamic range of the efficiency measurement. The error bars on each point denote the resulting 1σ confidence intervals for the coupling efficiency. Only counts from APD0 and APD3 were used, as their dark count rates were significantly lower than APD1 and APD2.

to map the entanglement to another ion species before measurement in combination with spectral filtering of the respective transitions on the optical link.

For this demonstration, however, the ions were disconnected from the optical link by displacing them from the narrow focus of the high-numerical-aperture imaging system. This had the advantage of being technically simple to implement – compared to e.g. free-space shutters, which might introduce mechanical vibrations, or might suffer from duty-cycle limitations –, as the apparatus had already been designed with shuttling-based experiments in mind. While the ions can be shuttled by a large distance ($\approx 1 \, \mathrm{mm}$) in the *HOA2* traps used for this work by sequential modification of the DC electrode potentials along the trap axis, only the the four second-nearest-neighbour electrodes (Q15, Q16, Q23 and Q24; see fig. 3.3) were updated for technical simplicity, as the temporary DAC system used to set the electrodes voltages at the time did not support high-rate, simultaneous updates for many channels.

As discussed in §7.4.1, after a herald event, a π pulse on the 674 nm $\left|S_{1/2}, m_J = 1/2\right\rangle \rightarrow \left|D_{5/2}, m_J = -3/2\right\rangle$ transition first maps the information to the slightly less magnetically-
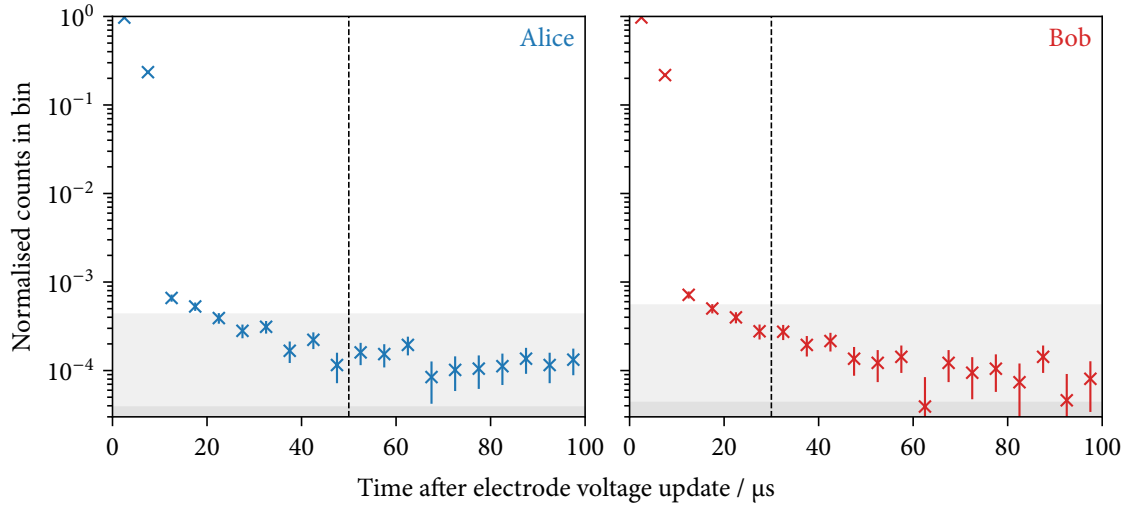
Figure 7.20: Time-dependent attenuation of photons coupled into the fibre link when shuttling ions away from the focus of the high-NA fibre coupling optics. The shuttling sequence was executed $\sim 3 \times 10^5$ times with the Doppler cooling lasers turned on, and the timestamps of heralding station APDs relative to the start of the sequence (the electrode DAC voltage update) were recorded. The plots show the count totals in 5 μs bins, relative to the count rate of an ion in the focus, with the error bars giving 1σ confidence intervals assuming Poissonian statistics for signal and background. The upper horizontal line up to which the background is shaded corresponds to the counts from the dark count rate of the used detectors, which were subtracted from the data. The lower line up to which the background is shaded more intensely represents the uncertainty in the background count rate measurement (assuming Poissonian statistics), which limits the dynamic range of the attenuation measurement. The dashed line at 55 μs (30 μs) marks the start of the analysis pulse and subsequent readout.

sensitive optical qubit transition. At this point, the link is still connected. Before applying the measurement-basis-dependent pulses, the ion is then moved 3 μm from the focus of the high-NA lens to greatly reduce its coupling to the link (see fig. 7.19). To separately study the dynamics of this process, we initialise the ion in a bright state, and shuttle it to the side many times with the fluorescence lasers turned on, monitoring the arrival times of photons at the heralding station. The resulting time traces are shown in fig. 7.20. After a delay of 30 μs to allow the low-pass-filtered electrode voltages to settle, Bob sends the round type signal (T_i) to Alice. This is the point from which isolation is necessary; here, the coupling has been reduced to $\sim 10^{-4}$.

Due to the large difference in wavelength, the probability for the 1092 nm photons generated from occasional $P_{1/2} \rightarrow S_{1/2}$ decays during state readout to couple into the single-photon-carrying fibre is so small that it would be very challenging to measure in our system (certainly smaller than the residual coupling for 422 nm photons, which make up $\approx$ 95 % of scattered photons). No photons from the 674 nm laser beam used to implement the basis selection could be detected at the output of the single-mode fibre either[11]. The coupling strength of the fibre mode to the ion for other wavelengths does not need to be

---

[11]While the 674 nm laser beams are quite strong (at the $\sim$ 1 mW level), the absence of measurable scatter on the 422 nm single-photon detectors is not particularly surprising given the fact that the beams propagate at

Ion state		$k = 0$	$k = 1$	$k = 2$	$k = 3$	$k > 3$	$\langle k \rangle$
Alice	Bob						
dark	dark	141 506	3864	80	2	0	0.0277
dark	bright	656 728	19 770	378	6	0	0.0304
bright	dark	670 039	18 685	322	7	0	0.0281
bright	bright	141 707	4347	76	3	0	0.0308

Table 7.5: Leakage after shuttling-based link isolation during the DIQKD link characterisation run (§7.4.2). The columns give the numbers of rounds during which k photons were registered at the central heralding station detectors and the respective means $\langle k \rangle$, conditioned on the classical outcomes measured on both sides. The background counts measured for the dark classical outcome 01 are, to within the statistical uncertainty, entirely due to detector dark counts.

considered: even if an attacker were able to e.g. create appreciable amounts of 674 nm or 854 nm radiation intensity at the ion by coupling a strong laser beam back into the fibre, leakage of information *into* the nodes does not need to be excluded, as its absence is ensured via the Bell test rounds.

Appropriate isolation of the nodes is an assumption in DIQKD that cannot be verified using device-independent techniques. If there is any doubt as to the amount of information leakage, isolation far beyond the required level should thus preferentially be established through independent technical means (mechanical shutters/…, as mentioned above), rather than trying to model the device behaviour in detail. Nevertheless, it is possible in principle to tolerate a small amount of residual leakage by shortening the length of the extracted secret key accordingly. We chose not to rigorously pursue this here, but in light of the somewhat limited extinction indicated by the above measurements (due to the small transport distance), it is reassuring that taking these effects into account would not compromise the claim of a positive key length. To verify the amount of shuttling-based isolation, we additionally monitor the heralding station detectors while the nodes read out the qubit states with the link disconnected during the link performance characterisation run. The count totals listed in table 7.5, from which we estimate the probability of a photon leaking from Alice during a bright measurement (classical outcome 0) as 6×10^{-4}, that from Bob (classical outcome 1) as 4.4×10^{-3}. Here, we factored in an additional ~ 25 % entangler coupling loss and ~ 70 % quantum efficiency of the single-photon detectors used, as well as their dark count rate, which accounts for most of the background – an ideal adversary would not be restricted by those. As the readout pulses take up almost the complete sensitive time window and the laser parameters are optimised to obtain a high count rate from the ions, extra laser light coupled back into the setup by an attacker could not appreciably increase the amount of information leaked. For the sake of this estimate, we can lump in the heralding station with one of the nodes (due to the observed asymmetry in extinction,

---

right angles to the collection axis, the large spectral separation, and the lack of immediate scattering surfaces along that axis given the slot beneath the trap axis in the *HOA2* traps.
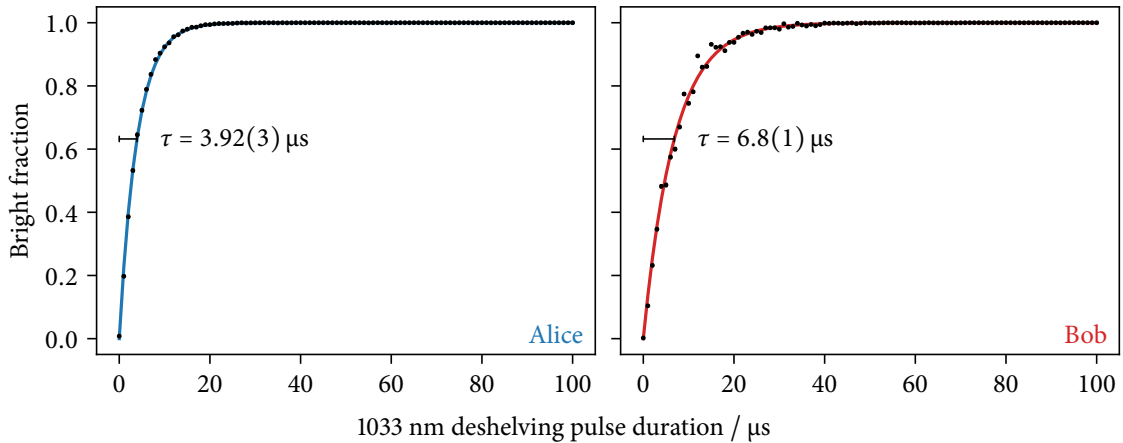
Figure 7.21: Population in the $S_{1/2}$ ground level after initially preparing the ions in the $\left|D_{5/2}, m_J = -3/2\right\rangle$ state and applying a de-shelving pulse of various duration (left: Alice, right: Bob). Exponential fits and the resulting time constants are shown, which differ due to a difference in power and spot size at the respective ions, but are in either case much shorter than the 100 µs after which the photonic link is reconnected by shuttling the ion back to the high-NA lens focus. Thus, the ion has entered the $S_{1/2} \leftrightarrow P_{1/2} \leftrightarrow D_{3/2}$ cooling cycle with very high probability.

preferentially Bob). This effectively eliminates the leakage from one node, as the amount of coupling from one input of the entangler apparatus back out of the other input is very small. As the probability for a bright measurement on Alice's side is 1/2, the probability of a photon leaking from either system during the measurement is 3×10^{-4}. The information leaked by this channel is upper-bounded by the binary entropy $h(3 \times 10^{-4}) = 4 \times 10^{-4}$, which is significantly lower than the key rate of 0.064 obtained in the DIQKD demonstration run, and thus does not impact the claim of a positive key rate.

Given the importance of the isolation assumption for the security of DIQKD, it might still be preferable to place a bound low enough to be entirely negligible on the leakage by an independent, macroscopically verifiable mechanism, such as through the addition of fast mechanical shutters.

7.6.3 Qubit state scrambling before link reconnection

Given the quantum-non-demolition nature of fluorescence-based state detection in trapped ions, Alice and Bob need to explicitly scramble the state of their qubits at the end of one measurement round, that is, before reconnecting the quantum link in preparation for the next entanglement generation attempt. Otherwise, the attacker might be able to shine in the beams required for state detection through the photonic link and gain information about the state before it is finally scrambled as part of the heralded entanglement attempt cycle.

Like the isolation process discussed in the last section, this scrambling is again explicitly a device-dependent process, as we rely on the knowledge that the applied reset pro-

cedure actually removes any trace of the measurement outcome from the system. This is, in fact, a concrete instance of a more general challenge to device-independent protocols, that of *memory* or *device reuse* attacks [BCK13]. Thankfully, dissipative processes that just lose the quantum information tend to be much easier to plausibly and robustly implement than coherent manipulation or measurements, especially for a qubit system of as low a complexity as a single atomic ion. In practice, Alice and Bob achieve this by simply applying the Doppler cooling lasers for 100 μs, including the 1033 nm laser to return population from the $D_{5/2}$ shelf level (that corresponds to a "dark" measurement outcome) to the cooling cycle. This is long enough for > 10 deshelving time constants to elapse, and many 422 nm cooling photons to be scattered. When the link is finally reconnected and the entanglement generation cycle resumes, the ion state is thus, for all practical purposes, already uncorrelated to the readout result.

In the long term, the reliance on a necessarily device-dependent reset mechanism could be reduced if quantum registers large enough to hold the complete set of qubits needed for the execution of the protocol are available, such that no reset would be needed within the protocol execution. In a realistic scenario where Alice and Bob wish to continuously generate fresh key material, the same issue would still arise when the apparatus is re-used for the next run of the protocol, however [BCK13].

7.6.4   *Sharing of laser sources*

For this demonstration, the two nodes are located in the same laboratory; it is thus convenient – and research-budget-friendly – to derive all the required laser beams from the same set of sources. This is, of course, not representative of a real-world deployment of DIQKD, where it would likely not be economical to route several additional laser signals between nodes. There could be two possible consequences of this for the interpretation of this proof-of-concept demonstration. First, it could completely render the device-independence analysis invalid, which would put the validity of the entire demonstration into question. Secondly, it could have consequences for the feasibility of a real-world deployment, if this tacitly led to the assumption of optical phase coherence between the nodes. However, neither of these cases applies to our experiment, as we will discuss in the following.

First, in practice, the laser beams are distributed between the systems by aid of a beam-splitter, which, when acting on the large coherent state of laser light, produces two independent coherent states in its output, linked only by the same coherent state amplitude parameter. Resorting to this argument, however, is not actually necessary to show that the sharing of the laser source is benign.

Rather, the only requirements placed on the measurement settings by the security proof are that the quantum state in each round is independent of the measurement settings, and that the measurements act in a local fashion on Alice's and Bob's subsystems (i.e. can be

described as a tensor product). The first is guaranteed by the assumption of trusted randomness, which is not revealed until the links are disconnected. The second follows from the fact that the two nodes occupy well-defined, disjoint regions in the laboratory (in principle, there would not be anything to stop us from putting the two halves of the setup into separate Faraday cages). Thus, it would not be an issue even if e.g. the shared 674 nm laser was actively controlled by an adversary; in fact, one of the advantages of DIQKD is precisely that it guards against coherent attacks like this.

In this picture, we neglect the potential leakage of information back through the laser supplies, as the coupling of each laser beam to the ion is very weak in absolute terms, and could be further suppressed using optical isolators. We do not attempt to quantify this further, as in the case that this was a concern, it would be much simpler to equip each node with its own laser sources instead, such that only the classical network and quantum links need to leave the nodes.

Lastly, using the same lasers for both nodes in our setup is really only beneficial in terms of equipment cost and laboratory space needed; the only amount of technological overhead avoided is that of occasional maintenance of the second set of lasers, as we already keep the rest of the control systems completely separate. In an application with a larger distance between nodes, a separate set of lasers could easily be used instead, as we do not rely on any optical phase coherence between the nodes. The only synchronisation requirement is for the entanglement-carrying photons to be well-matched in frequency and arrival time at the heralding station compared to their 7 ns duration. In particular, we do not require the phases of the 422 nm picosecond pulses or the absolute phase of the coherent 674 nm laser pulses to be matched between the nodes or be stable. The relative phase between projective measurements at each node is determined solely by the 14 MHz RF beat note between the two 674 nm pulses. As such, the most stringent synchronisation requirement derives from the time overlap of the 7 ns photons at the heralding station, which is not particularly challenging to achieve with modern time dissemination techniques – in the simplest case, a clock signal could be directly sent from one node to the other.

As a side note, the apparatus used for the first demonstration of device-independent randomness extraction [Pir+10] also employed a shared set of laser sources, although the authors do not appear to explicitly address this – it is indeed irrelevant for the device-independent analysis, as long as the measurement settings are chosen locally. More fundamentally, some amount of sharing of optical or electronic equipment is always going to be necessary in an apparatus based on the measurement of remote entanglement, as, at the very least, a common phase reference for the qubit measurements needs to be established. In addition to this, both platforms based purely on photonic technologies and the heralded entanglement generation schemes employed for atoms, ions or NV centres – including the setups used in loophole-free Bell tests – also require the fine alignment of single-photon time slots. Whichever one of these timescales happens to be more demanding for a given

technology, at least some synchronisation equipment to ensure the time references are matched will thus be required.

7.6.5   Trusted randomness

As the security of DIQKD rests on the fact that an attacker cannot predict the measurement settings applied in each round ahead of time, it is desirable to use a source of randomness rooted in some well-understood hardware process, e.g. a quantum random number generator (QRNG). Note that our requirements are technically much less challenging than for the loophole-free Bell test experiments, where considerable engineering effort was spent on the basis selection circuitry: for DIQKD, the isolation assumption supersedes the requirement for space-like isolation, as discussed in §7.6.1, and we thus are not concerned with the "freshness" of the produced random bits. Without any latency requirement, we can make use of the trusted local classical computation capabilities to apply any "whitening"/randomness amplification algorithms as necessary, and do not need to integrate the random number generation with the real-time control system particularly tightly.

In fact, this is precisely the approach we take in our demonstration experiment. To illustrate a possible implementation, Alice and Bob use an ID Quantique *Quantis QRNG USB 4M* quantum random number generator module, which is based on beamsplitter photon counting statistics of a heavily attenuated light source. A random number string considerably longer than necessary for the experiments is generated ahead of time at each node and stored in the memory of the respective control PC. For each acquisition block, the round types and measurement settings are derived from these strings on the PCs and transmitted to the FPGA core devices.

To actually establish trust in the random nature of the device outputs, an analysis of the internal workings of the QRNG module would be necessary (which has been performed by various commercial certification laboratories, and recently by external researchers [Pet+22]). We did not perform such an analysis ourselves, but to exclude any major issues in the quality of the generated outputs, analysed a sample of $19\,121\,635\,328$ bits (far exceeding the $\sim 10^7$ bits consumed during the experiment) for a number of statistical defects using two widely used randomness test suites. *dieharder* [Bro20, version 3.31.1] with default settings[12] revealed no failures and 2 "weak" tests (out of 114 total, which is within statistical expectations). All tests from the "Rabbit" and "Alphabit" batteries from *TestU01* [LS07, version 1.2.3], run over the maximum of $2^{32} - 1$ bits, passed as well.

Curiously, despite these industry-standard tests passing, a simple manual tally of the number of zeros and ones in the full output did reveal a statistically unambiguous deviation from a perfectly uniform distribution. The sample contained only $9\,559\,817\,952$ zeros, which gives a probability point estimate of $0.499\,948$, or a two-sided p-value of $< 10^{-46}$

---

[12] `dieharder -a -g201 -f randomness.dat`, where "randomness.dat" is a file containing raw binary data.

for the null hypothesis of a uniform distribution. A small bias of this kind was also documented in ref. [HH20], and hints at a small design flaw in the on-module "unbiasing" stage. As there are no latency requirements in our experiment, the bias can easily be removed in post-processing. Using the universal-hashing-based extractor implemented in the ID Quantique software (with the manufacturer's default 2048×1972 matrix) left no statistically significant bias behind in a sample of $8 \times (2^{32} - 1)$ bits. Note that over the relatively small number of protocol rounds in our demonstration, the bias remains well within what would be likely just from statistical noise even without this extra post-processing step.

To avoid accidentally introducing any small bias between the protocol realisation and the assumptions in the security proof, we derive the classical inputs X and Y in a way that can very easily be seen to be correct, but is somewhat more wasteful in randomness than necessary. The nodes both draw directly from the QRNG output for the Bell test inputs $X, Y \subset \{0, 1\}$. We choose a power-of-two fraction for the test round probability ($\gamma = \frac{13}{2^8}$), and Bob implements the round type choice in the straightforward way by consuming eight random source bits. For key generation rounds, the random bits allocated to the Bell test basis choice are simply discarded. This could be optimised further in a scenario where random bits are slow to generate.

Ultimately, though, note that the above considerations serve simply to make the correct implementation of the random choices more plausible. The availability of trusted randomness remains an assumption in DIQKD, as discussed in §7.2.1, and is not addressed by our experiment. In a real-world high-stakes deployment, this would likely be addressed through a scrupulous engineering analysis of the random number generator hardware. Sufficiently paranoid users could also make use of device-independent randomness amplification and multi-source extraction techniques, as e.g. described in ref. [KA20], to combine randomness from multiple dissimilar sources.

7.7   Discussion

The results in this chapter mark the first demonstration of device-independent quantum key distribution, where the used quantum devices are self-tested as part of the protocol execution. Crucially, we have successfully obtained a key in a complete end-to-end demonstration: by combining our state-of-the-art experimental platform for remote entanglement with a number of improvements to the theoretical analysis and the classical post-processing pipelines, we were able to execute the detailed protocol on separate nodes from start (i.e. fixing the protocol parameters) to finish, where the nodes each have an identical copy of a concrete secret key. This key is secure under the most general assumptions, where the adversary can act arbitrarily on the quantum state (e.g. including entanglement across rounds), and the measurement devices are allowed to act in a fashion correlated to

the outcomes of previous rounds. This unambiguously demonstrates that provably secure, device-independent quantum cryptography is possible with real-world devices.

Having cleared this straightforward criterion for success, the generation of a key longer than any key material consumed, is particularly reassuring, as DIQKD relies on many technical details from experimental quantum information processing as well as classical and quantum information theory. The merits of a claimed demonstration of DIQKD where the effort stopped short of obtaining a key would be difficult to evaluate for readers who are not domain experts in both the relevant experimental and theoretical disciplines, as it would be hard to judge which pieces are actually missing, and how feasible it would be to fill those gaps. Conversely, as long as the entropy accounting in the security analysis is done correctly, obtaining a key directly demonstrates that there are no unexpected surprises in any of the many components, neither in the stability of the experimental apparatus, nor any potentially unrealistic assumptions as to e.g. the performance of the error correction code. In view of this, the detailed description of the protocol and state-of-the-art performance of the post-processing pipeline used in our work should hopefully be of help to future experimental efforts on other platforms as well.

7.7.1 Related experimental work

After publication of our initial manuscript as a preprint [Nad+22], two other manuscripts appeared which make claims related to an experimental DIQKD demonstration. Neither of them, however, attempted to actually obtain a secret key, nor could they have easily succeeded in doing so, as we shall briefly discuss.

W.-Z. Liu et al., *Toward a Photonic Demonstration of Device-Independent Quantum Key Distribution* (Hefei/Shanghai). The authors of ref. [Liu+22] pursue a fully photonic approach, where entanglement between 1560 nm photons is created through spontaneous parametric downconversion at a central station and directly measured at Alice's and Bob's nodes. As described in §7.3.2, the lack of heralding is a challenge for purely photonic implementations, but if it can be surmounted, high raw key rates are potentially available in a technically simple fashion. To this end, the authors present a more complex DIQKD protocol, where a novel combination of randomised post-processing and careful post-selection increases the tolerance to finite detection efficiency. Through careful optimisation of the detectors and sources, an overall system efficiency (including both entangled state generation and detection) of 87.5 % is reached, resulting in a detection-loophole-free Bell inequality violation of $S = 2.0472$.

No attempt to extract a key is made, but the authors claim that their performance characterisation is equivalent to an asymptotic secret key generation rate of 466 bits/s over 20 m of fibre, or 2.6 bits/s over 200 m. The setup as presented, however, does not include any

fast polarisation modulators to select the polarisation basis in real-time (for the performance analysis, wave plates are rotated mechanically to alternate between basis settings in blocks). As the system efficiency is already very close to the threshold of 86.2 % quoted by the authors where the key rate vanishes entirely even in the asymptotic limit, fitting one modulator per node into the efficiency budget remains as a significant challenge; even loss-optimised free-space electro-optical modulators typically exhibit losses of at least 1–2 %. Furthermore, the security analysis is only performed in the collective (i.i.d.) limit; as the protocol relies on post-selection, it is not obvious how the key rate will be affected when transitioning to the full setting of coherent attacks.

W. Zhang et al., *A device-independent quantum key distribution system for distant users* **(München).** In ref. [Zha+22], the authors describe a platform based on single optically trapped ^{87}Rb atoms. Impressively, this experiment was carried out with the nodes located in two different laboratories separated by 400 m (700 m of fibre), using photons at 780 nm for the quantum link as previously employed for a loophole-free Bell test experiment [Ros+17]. After each herald, the atom state is detected via a destructive state-selective photoionisation technique, which lends itself to making later classical leakage of the internal state implausible. The authors employ a protocol broadly similar in structure to the one described here, but with an extra measurement setting to increase tolerance to larger quantum bit error rates [Sch+21b]. Measuring 3342 Bell pairs over 75 h, they show that their link performance very likely lies in the region where a key could be extracted in the asymptotic limit of infinitely many protocol rounds.

No attempt at practically implementing the DIQKD protocol is made, however. Presumably this is because it was experimentally infeasible to continuously acquire data for the extremely long durations that would be required given the modest entanglement quality ($S = 2.58(8)$, $Q = 7.8(9)$ %) and low rate of 0.012 s^{-1}: depending on the choice of security parameter, the experiment would have to be run continuously for on the order of one year to obtain any key at all, even in the case analysed by the authors where only collective (i.i.d.) attacks are considered, and several dozen years to reach security against general attacks with parameters similar to our result.

7.7.2   *Possible extensions*

Leaving aside the possibility of incremental improvements to the quantum link performance – while work on e.g. the entanglement generation rate is important with a view towards real-world deployments, it wouldn't lead to any results conceptually different from our existing DIQKD result –, two avenues of investigation for future proof-of-principle experiments should be highlighted.

The first concerns the assumptions made in the protocol, and their experimental plausibility. For instance, as discussed in §7.6.3, future large-scale quantum memories could

make it possible to execute all the measurements required for the protocol in one go after accumulating all the required entanglement, perhaps making it easier to ensure isolation of the quantum link during that one measurement period. With experimentally achievable Bell inequality violations increasing, it might also be interesting to revisit protocols which do not assume the validity of quantum theory [BHK05], though it is unclear how node isolation would then be ensured.

The second direction concerns the validation of improved DIQKD protocols, and the security analyses thereof. One trivial improvement to the protocol presented here is to swap the direction of the round type announcements T_i, such that no communication from Bob to Alice needs to be authenticated. For geographically distributed applications, where speed-of-light delays in communication can be significant, protocols that avoid this need for real-time signalling without requiring sifting could be of interest [Tan+20]. There are quite a few proposals for other improvements to the certifiable key rates in certain regimes of device parameters, for instance through more general measurements [Sek+21], the addition of small amounts of trusted noise [Ho+20], or two-way reconciliation schemes [TLR20]. A systematic comparison of such proposals, for some of which security against coherent, non-i.i.d. attacks has not been shown yet, would be useful for future practical demonstrations.

7.7.3   Outlook: DIQKD for real-world security

While our experiment already clears the bar for a demonstration of DIQKD in a realistic setting – we can obtain finite-sized keys in reasonable amounts of time, and the nodes could easily be separated completely by purchasing a second set of laser sources and moving one of the ion traps to a different laboratory –, significant work still remains to push it closer to relevance for practical quantum key distribution settings. As DIQKD is agnostic to the system details, any further incremental improvements to entanglement fidelity, entanglement generation rate, transmission distance, or compactness of technical integration of the nodes will also incrementally improve the DIQKD performance in the same way, given that the protocol assumptions remain fulfilled. We thus leave a discussion of these aspects for chapter 8, where we discuss them also for other, more general applications.

Taking a more conceptual, longer-term view, with DIQKD having long (and rightly) been discussed under headlines such as "the ultimate physical limit of privacy" [ER14], does this mean we have now solved communication security, achieved "unhackable" security? Even leaving questions of practicability aside, every reputable practitioner would answer in the negative, as such a question is founded in a fundamental misunderstanding of the challenges in securing real-world devices.

One way to dispel such myths is by comparison to the classical case. After all, the only thing keeping various classical cryptosystems from obtaining the same "unhackable" label is their reliance on the bounded computational capabilities of any attacker. But even

though the latter is still an excellent assumption for most widely deployed algorithms – at least until a major breakthrough in cryptanalysis or large-scale fault-tolerant quantum computing –, computers are often hacked and supposedly secure communication systems are frequently broken. Rather than through attacks on their mathematical foundations, this is due to implementation mistakes and incidental vulnerabilities, such as those to side-channel attacks. Any DIQKD implementation can suffer from these issues as well, in the classical part of the protocol, but not least also in the additional hardware now necessary for the quantum part.

This illustrates that to discuss implications for practical, real-world security, DIQKD has to be viewed through the lens of security engineering – a careful analysis of attack surfaces and various trade-offs in system complexity and performance. QKD uniquely provides information-theoretic security, even in the face of an adversary with unlimited computational power, at the cost of requiring a quantum link. But "conventional" QKD also adds the trusted models for the quantum device behaviour to the attack surface. DIQKD allows one to avoid this increase in attack surface, removing assumption (vi) from the list discussed in §7.2.1. In fact, the remaining assumptions (i)–(v) could be considered the bare minimum necessary for a meaningful encryption system, and we prove security in a fashion that is composable with other cryptographic primitives and algorithms [PR22]. As such, while DIQKD can never guarantee complete "black-box" security in itself – to reprise the discussion from §7.6.1, if the classical devices used contain a covert transmitter relaying the key material to the outside, no security can ever be achieved –, device-independent protocols can offer practical engineering benefits over classical and other QKD systems.

We hope that our demonstration of a system close to the ultimate theoretical limit of security can serve as a valuable example in a discussion about the level of security actually achievable in practice, illustrating for instance the practical challenges in ensuring isolation between rounds, the potential role of memory attacks, and similar issues.

8   *Outlook*

In this thesis, I have presented the theoretical study and experimental characterisation of a photonic link between distant ^{88}Sr$^+$ qubits in two separate ion traps. This was achieved through the collection of photons from spontaneous decay between levels with total angular momentum 1/2. This novel collection geometry is advantageous from a technological point of view; the symmetry properties of single-mode optical fibres prevent the loss of polarisation purity. The fidelity and rate reached currently represent the state of the art in remote entanglement generation, which allowed us to demonstrate two quantum networking tasks; one cryptographic task with stringent go/no-go requirements on the experimental performance (chapter 7 and ref. [Nad+22]), and another in using entanglement to enhance the frequency comparison of remote atomic clocks [Nic+22]. A discussion of more specific conclusions, including comparisons to prior work in the respective domains, can be found in §§5.7, 6.4, and 7.7. Here, concluding this thesis, I more generally discuss avenues for further improvements of the photonic remote entanglement generation link, which would increase the performance in the elementary two-qubit applications already demonstrated, and make more complex, multi-qubit distributed quantum information protocols feasible to realise.

8.1   *Potential for increased link performance*

8.1.1   *Increasing the entanglement generation rate*

Without any major changes to the apparatus, the entanglement rate could still be considerably increased from the $184\,\mathrm{s}^{-1}$ observed thus far. Using superconducting nanowire detectors [Esm+21; NTH12; Val+14; Sli+16], for which detection efficiencies in excess of 95 % have been demonstrated, in place of the APD detectors, and further optimising the Bell state analyser fibre couplings, a factor of 2.5× in two-photon coincidence rate should be feasible. The coupling efficiency of the ion-fibre interfaces is still below expectations; improving it to 85 % of the theoretical value for NA 0.6 would lead to an additional 2.9× increase in rate. Finally, with a different driver for the picosecond laser pulse picker, and some optimisation of the system latencies (perhaps adding some electro-optic modulators for the state preparation beams as well), the attempt rate could feasibly be increased to

$\sim 4\,\text{MHz}$. All in all, this would amount to a $\sim 30\times$ increase in rate, reaching an entangled state generation rate of $\sim 5\,\text{ms}^{-1}$.

At this point, the average time for the generation of a remote entangled ion pair would be comparable to the duration of typical laser-driven local two-qubit entangling gates. For modular quantum information processing applications (or if a high attempt rate is not feasible), a further increase in collection efficiency would still be useful. To do this in free space, one improvement ($\sim 25\,\%$) would be to better match the radial profiles of the (approximately Gaussian) fibre mode and the (more top-hat-like) imaged ion emission. This could be achieved with standard imaging optics through the use of fibres with a tailored mode profile or with matching optics integrated onto the fibre face [MF11; Gis+16; GVB16]. In general, adaptive optics might help improve the coupling efficiency, but would significantly increase the interface complexity. As discussed in chapter 4, an increase in numerical aperture in theory does not affect the ion–photon entanglement fidelity, even when collecting from an entire half-space. Various "soft" engineering constraints limit the numerical aperture achievable in practice (in particular, parts of the optical system might need to be moved inside the vacuum system), but NA 0.8 seems feasible. The collection angle could be further increased through the use of mirrors, whether macroscopic in nature or integrated with the trap electrode structure [Shu+10; Fis+14; Gha+16]. If the mirror spans close to 2π solid angle, a further improvement of collection efficiency should be possible through destructive interference from a quarter-wavelength step in the mirror profile [Ara+20].

It would also be possible to increase the entanglement generation rate by adopting protocols based on a single-photon herald signal. The interferometric stability requirement for the photonic links could be addressed using protocols which incorporate a stage of entanglement distillation tailored for this case ("extreme photon loss" [CB08]). Ion motion might preclude reaching high fidelity levels unless carefully controlled, however [Slo+13], and the ground-state cooling operations possibly required could significantly affect the achievable repetition rate. Whether such a scheme could ultimately increase the system performance over two-photon protocols will require more detailed investigation.

8.1.2   Increasing the distance

In the current demonstration, the photons propagate through $\sim 4\,\text{m}$ of single-mode fibre. This could be extended to building-scale distances, such as relevant for large-scale modular quantum computing. For longer distances, however, the loss in conventional step-index silica glass fibres represents a major challenge ($\sim 30\,\text{dB/km}$, i.e. halving the rate every $100\,\text{m}$). For an ideal fibre, the main loss mechanism is Rayleigh scattering ($\propto \lambda^{-4}$); the optimum is found around $\approx 1550\,\text{nm}$ (known as the "C band" in telecommunications), where losses as low as $\sim 0.14\,\text{dB/km}$ have been achieved in a research setting [Tam+18] (for longer wavelengths, losses from coupling to molecular vibration dominate). Optical

difference frequency generation using non-linear materials can be operated in the single-photon regime ("quantum frequency conversion") while preserving the coherence of the ion–photon entangled states, allowing the C band to be accessed even though none of the laser-cooled ion species directly have a transition in this wavelength region. In trapped ions, entanglement between a ^{40}Ca$^+$ ion in a cavity and a 1550 nm photon after transmission through 50 km of fibre has been demonstrated [Kru+17; Kru+19]; using neutral rubidium atoms recently also the "full" heralded entanglement generation process between two nodes across 33 km of fibre using down-conversion to 1517 nm [vLee+22].

In such frequency conversion systems, it is technically advantageous if the pump laser has the longest among the three relevant wavelengths, as this avoids background from non-phase-matched broadband spontaneous parametric downconversion and spontaneous Raman scattering [Pel+11], which would lead to spurious heralds and thus decrease the remote entanglement fidelity. For source (atom) wavelengths shorter than ~ 780 nm, which cannot be operated in this regime, either more aggressive spectral filtering is necessary, or a two-stage approach is necessary, both of which will have a certain efficiency (rate) overhead. Direct conversion of the 422 nm wavelength to the C band was demonstrated[1] [Wri+18], but the conversion efficiency of 1 % would benefit from optimisation, and further filtering would need to be applied to obtain reasonable fidelities given the only ~ 4 % source photon collection probability in current experiments. Two-stage conversion from ^{138}Ba$^+$ photons (493 nm) to 1534 nm has also been demonstrated [Han+21], though at an efficiency of only 1–5 %.

In a realistic, spatially separated deployment of quantum frequency conversion, the pump lasers would need to be stabilised with respect to each other. As the relative linewidth merely needs to be small relative to the photon bandwidth, this is relatively easy to achieve (e.g. through distribution of seed light alongside the quantum link, or by referencing the lasers to a frequency comb). If polarisation encoding is used, the birefringence would likely need to be actively stabilised for longer, deployed runs of fibre, but this is routinely achieved in practice through periodic calibrations (see e.g. ref. [Tre+09]).

In the case of ^{88}Sr$^+$, transmission losses at the 1033 nm and 1092 nm wavelengths of the $P \leftrightarrow D$ transitions can be limited to < 0.7 dB/km in loss-optimised pure silica core fibres [Nag+02]. For medium distances, direct transmission of such photons might thus be a viable, technically simpler alternative to quantum frequency conversion. The few-percent branching fraction from the P to the D levels would place a significant limit on the achievable system efficiencies, though, if not enhanced through the use of optical cavities.

As an alternative to step-index silica fibres, hollow-core fibres, where the mode is almost entirely guided in a void between resonant elements, have the potential of vastly reducing losses from Rayleigh scattering. Low-loss hollow-core fibres, particularly for shorter wavelengths, are still an active area of research, but early results appear promising

---

[1] For technical reasons, the wavelength in this demonstration was actually 425 nm.

for reaching few dB/km losses also at 422 nm (see e.g. ref. [Sak+20] for a demonstration at 600 nm). Free-space transmission, which has been used for QKD between distant ground stations [Sch+07], as well as between ground stations and a satellite [Lia+17], could also drastically reduce the link losses across long distances.

8.1.3   Optical cavities

Some of the earliest proposals for the realisation of light–matter interfaces for quantum networking have involved optical cavities (e.g. [Cir+97]), and cavities have prominently been used to link optically trapped neutral atom qubits with some success (though still limited in fidelity/entanglement rate; see the literature comparison in chapter 6). Optical cavities have been investigated in conjunction with ions in radio-frequency traps as well, where they promise an enhanced collection efficiency for a given solid angle by modifying the optical mode structure at the ion position, limited in principle only by the residual losses in the cavity mirrors. This way, the weaker $P \leftrightarrow D$ infrared transitions in earth-alkaline ions could also be directly accessed. There has been considerable experimental progress in recent years (see ref. [Kel22] for a recent introduction and review). For instance, ion–photon entanglement with an overall collection and detection efficiency of 46.2(3) % was recently demonstrated using $^{40}Ca^+$ [Sch+21a].

While clearly of interest for high-performance photonic links between trapped-ion qubits, optical cavities are not without their challenges in this context as well. Firstly, and perhaps least importantly, the photons generated through optical cavities tailored for efficient extraction tend to be several orders of magnitude longer in duration, limiting experimental repetition rates, and potentially increasing sensitivity to drifts. More importantly, electric field noise from and/or charge accumulation on the dielectric cavity mirrors, which are necessarily in close proximity to the ion for efficient designs with a small mode area at the ion, presents a major challenge in practice (see e.g. ref. [Tel+21]).

A further complication for the use of cavities in the generation of heralded remote entanglement is that the photons from each node need to be indistinguishable to a high degree in order to obtain good fidelities. This requires a variety of cavity and driving parameters to be matched between the nodes; the observation of high-fidelity ion–photon entanglement in a single-node experiment thus does not necessarily translate to the heralded remote entanglement case. For instance, in a recent demonstration of remote entanglement between cavity-enhanced ion trap nodes, differential fluctuations in the cavity lengths were a major contributor to the overall fidelity error, despite active stabilisation of the same [Kru+19]. If decay channels of differing polarisations are used, the matching requirement includes the cavity birefringence; as the manufacturing of cavities with sufficiently low and well-controlled bireringence is technically challenging, this might motivate use of other (e.g. time-bin) qubit encodings [Kas+23]. Indistinguishability is much

easier to ensure in free-space spontaneous emission, where the qubits are, in the absence of modulation from motion or magnetic fields, guaranteed to be identical.

8.1.4   Increasing the Bell-state fidelity

As shown in chapter 4, fidelities in excess of 99 % should in principle be achievable even under a realistic choice of parameters for a wide range of experimental imperfections. At present, the errors observed in the two-node network appear to be dominated by residual polarisation mixing, presumably due to aberrations in in the imaging optics, or non-uniform birefringence e.g. in the vacuum viewport. Further, quantitative investigation is not only necessary to confirm this hypothesis, but might also lead to further increases in collection efficiency, if the cause is in some (non-symmetric) aberrations in the imaging system.

At the $\mathcal{F} > 95\,\%$ level, however, a whole number of effects are potentially relevant. In the face of longer distances, for instance, heralding detector dark counts or, if quantum frequency conversion is employed, downconversion noise, can quickly become non-negligible, and the achievable performance should be carefully modelled. To consistently reach even higher fidelity levels, it might be necessary or preferable to use the high-fidelity local entangling gates available in trapped ions to probabilistically combine multiple raw Bell pairs into one of higher fidelity (see next section).

8.2   Towards more complex networks and network nodes

Trapped ions are attractive as a platform for complex quantum networking applications not only because their identical, well-controlled nature lends itself to high-fidelity remote entanglement generation as demonstrated in this work, but importantly also because of the high coherence times and the well-developed toolbox of high-fidelity operations available in trapped ions (state preparation, readout, universal gate sets). Here, trapped ions generally also offer state-of-the-art performance across all qubit technologies; for instance, the Oxford group has previously demonstrated two-qubit entangling gates reaching a $^{43}\mathrm{Ca}^+$–$^{43}\mathrm{Ca}^+$ Bell state fidelity of 99.9(1) % [Bal+16], and a $^{88}\mathrm{Sr}^+$–$^{43}\mathrm{Ca}^+$ mixed-species Bell state fidelity of 99.7 % [Hug+20]. Our ion-trap quantum network nodes support using the $^{43}\mathrm{Ca}^+$ species[2] in addition to the $^{88}\mathrm{Sr}^+$ interface ions, and are based on microfabricated surface traps. This makes plausible the claim that the photonic interface is compatible with a technology at least in principle scalable to a large number of ions per node, and complex interactions between them.

---

[2]Control of extended ion crystals and mixed-species gates have already been implemented, though at the time of writing, the fidelity of the $^{88}\mathrm{Sr}^+$–$^{43}\mathrm{Ca}^+$ gates is still limited by motional coherence issues.
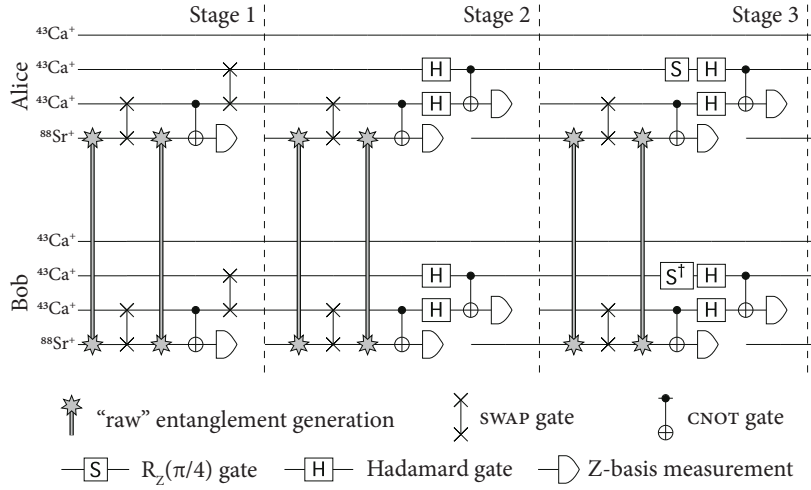
Figure 8.1: Circuit representation of a cascaded entanglement distillation protocol between two identical ion trap modules "Alice" and "Bob". Two "raw" Bell pairs generated over the photonic link are used in each of the three stages, reducing the infidelity from ϵ for the raw Bell pairs to $\approx \frac{2}{9}\epsilon^2$ after success of the entire protocol. Depending on the outcome of the parity measurements, the respective stage or the entire process might need to be restarted. The two inactive ^{43}Ca$^+$ "application qubits" hint at the remainder of the computation, where the distilled entanglement can subsequently be used. For instance, to realise a remote two-qubit gate Alice and Bob would both apply an extra two-qubit gate between the distillation result and application qubit, measure the former, and perform any necessary local unitary corrections. (Figure adapted from ref. [Nig+16].)

More practically, this design was in particular chosen with the aim to support the realisation of universal entanglement distillation protocols to generate remote entanglement at or near the fidelity level of local operations. In particular, our system can support the three-stage distillation protocol described in ref. [Nig+16] (see fig. 8.1), using five ions per trap (three ^{43}Ca$^+$ for storage, one ^{88}Sr$^+$ as photonic interface qubit, another ^{88}Sr$^+$ as sympathetic coolant), and relying on (few) crystal configuration changes (shuttling/splitting/merging, without requiring rotations) to choose different subsets of ions to interact instead of tightly-focussed laser beams. Given the remote entanglement fidelities achieved here and local operations at the state of the art previously demonstrated in Oxford, simulations indicate that a fidelity of $\mathcal{F} \approx 99.8\%$ should be reachable this way using six Bell pairs. This would unambiguously be compatible with typical error-correcting codes considered for large-scale fault-tolerant quantum computation (though specialised codes taking into account the increased error rates but also the flexible topology for photonic links between modules will likely be advantageous for this). In the near term, given that the $\approx 4\%$ error in the remote ^{88}Sr$^+$–^{88}Sr$^+$ state observed here is not uniformly depolarising but contains considerable structure (see §6.3.1), only one distillation layer could already achieve $\mathcal{F} > 99\%$.[3] For this, one ^{88}Sr$^+$–^{43}Ca$^+$ ion pair per node is sufficient, used to probabilistically combine two Bell pairs with $\sim 60\%$ probability for success.

---

[3] Numerical calculations by B. Koczor and S. Benjamin (private communication), performed with the assumption of perfect node-local operations (including the ^{43}Ca$^+$–^{88}Sr$^+$ entangling gates).

Mixed-species operation of the Alice and Bob nodes would also enable the prototyping of extensions to our recent entanglement-enhanced frequency metrology demonstration [Nic+22] to a setting more relevant to state-of-the-art atomic clocks using techniques akin to quantum logic spectroscopy [Sch+05], where one ion species well-suited to cooling/readout and remote entanglement generation would be combined with another with favourable properties for a frequency standard. ^{43}Ca$^+$ is not the most favoured species regarding systematic shifts, but should be perfectly adequate for the experimental investigation of such techniques, particularly as a first-order magnetic-field-insensitive 729 nm transition is readily accessible in our setup (at $B = 0.338$ mT and $B = 0.496$ mT [Ben+07]).

Another promising avenue of investigation would be to combine one such mixed-species ion trap node with possibly significant local quantum resources ("server") with a simple node that can just receive and measure photons in a variable basis ("client"). Interestingly, even such restricted receiving nodes are enough to enable algorithms not achievable between classical client and server computers [WEH18]. One particularly counter-intuitive capability is that of "blind" quantum computing [BKF09], by which it is possible for a client to have secure access to remote quantum computers, without the operator of the computer learning anything about the executed algorithms or processed data (see fig. 8.2).

Regarding the scaling to more complex devices, the current approach of fibre coupling using bulky, refractive high-NA optics naturally produces only one interface site on the ion trap chip, the demagnified image of the fibre tip. Multiple sites could be coupled to multiple fibres positioned in the right arrangement on the image sides, or to the same fibre through scanning optics (but with increased requirements on optical lens design and alignment). Even with just one interface site, multiple ions could still be connected to the optical link in turns by shuttling them around as necessary in a segmented electrode trap. Such a shuttling approach would also be compatible with optical cavities or exotic high-NA collection optics. For this, one could imagine a long racetrack-shaped trap section where ions are circulated in conveyor-belt fashion belts until heralding success is reached, at which point they are ejected into the rest of the quantum information processor chip (and returned after use). Using a large number of ions, an arbitrary amount of link latency between quantum network nodes could be tolerated while keeping a high attempt rate, in precisely the same way that buffering hides latency in classical computer science. This could be implemented in a pipelined architecture where the different operations (state preparation, excitation/spontaneous emission, waiting, re-cooling after failure) are performed in spatially different zones, perhaps in analogy to that developed for quantum logic gates by transport through a fixed laser beam [Lei+07; dCle+16]. In theory, one could reach entanglement attempt rates of many MHz like this, even over longer distances. The only limit would be due to the ~ 10 ns time scale of spontaneous emission typical for candidate optical transitions in ion species of interest (without cavity enhancement), where increasing the repetition rate further would necessarily reduce the probability of collecting any one

(a)

(b)

(c)

(d)
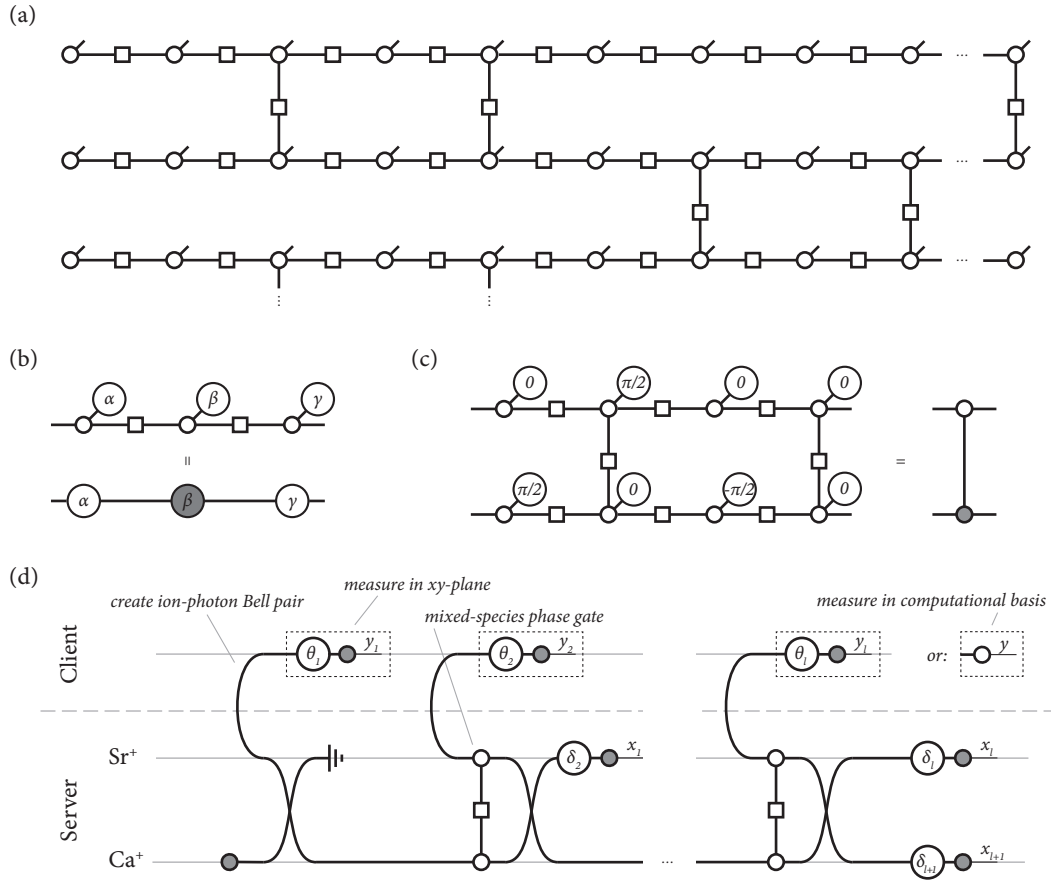
Figure 8.2: Universal blind quantum computation using a photonic interface to a trapped ion processor (illustrated in the graphical formalism of Coecke et al. [Coe10]): *(a)* The "brickwork" state [BKF10], universal as a resource state for measurement-based quantum computation with single-qubit xy-plane measurements (indicated by unfinished ticks). Qubits are initialised in $|+\rangle = |\varphi(0)\rangle$, $|\varphi(\alpha)\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle + e^{i\alpha}|1\rangle\right)$ and entangled using cz in a two-dimensional pattern as indicated. The protocol can be blinded if the client prepares the states in $|\varphi(\theta_k)\rangle$ instead (with the phases θ_k hidden from the server), who then carries out the rest of the computation. *(b)* Fragment of a brickwork state used to implement an arbitrary single-qubit rotation (Euler angles). *(c)* Fragment of a brickwork state used to implement a cnot gate. *(d)* Proposed experiment using a ^{88}Sr$^+$–^{43}Ca$^+$ crystal. The photon measurement angles θ_k defining the initial phase of the cluster state nodes are chosen by the client at random, the ion measurement angles δ_k communicated to the server are based on the former and feed-forward corrections from previous steps (which can be exchanged in real-time due to the long ^{43}Ca$^+$ coherence times, avoiding postselection). $\{\delta_k\}_k$ could be fixed if the $\{\theta_k\}_k$ are updated in real-time along the lines of ref. [MF13]. With the one ^{43}Ca$^+$ storage ion shown, arbitrary single-qubit and 1-deep two-qubit gate circuits can be implemented this way. Measuring the photon along the Z axis instead creates a "trap qubit" which splits the cluster state (as used in verifiable blind quantum computation).

photon.The technical challenges of shuttling ions quickly without significant motional excitation might also pose a limit in practice, though.

The integration of laser beam delivery into microfabricated ion trap chips using lithographically defined waveguides and free-space couplers has recently seen considerable attention [Meh+16; Meh+20; Moo+22]. Particularly for sensing or quantum networking applications deployed in the field, such integration could already be interesting even for relatively simple remote entanglement experiments. The use of on-chip waveguides could potentially be extended to single photon collection and routing. Across larger ion trap chips, optical connectivity could potentially offer greater connection rates than achievable using shuttling (including the necessary cooling steps), though polarisation purity, coupling efficiency, and coverage of the large solid angle necessary to obtain appreciable coupling efficiencies will potentially present a challenge.

Whether for the modular design of a large-scale quantum information processor, or remote communication or sensing applications, an extension to multiple quantum network nodes (as recently demonstrated for three nodes in NV centres [Pom+21]) will be of considerable interest. A direct connection to more than one node could be achieved by placing photon connection optics on more than one side of the trap centre (e.g. up to six along the cartesian coordinate axes), or more likely, through a single interface per site or node combined with an optical switch matrix to dynamically reconfigure the links to provide arbitrary long-range connectivity [DM10].

A Truncated multi-mode Fock space formalism for few-photon interference

When describing the heralding station in a remote ion–ion entanglement experiment (§4.4), we are interested in at most one photon from each ion, two in total, distributed between up to eight different spatial/polarisation modes. They are all (potentially) distinguishable, and the detectors time-resolved, necessitating the time coordinates to be accurately traced through the calculation as well. This makes it desirable to express these calculations in such a way as to be amenable to computer-assisted manipulations. For the analysis of photonic circuits in the quantum computing context, a number of advanced algebraic approaches exist, some of which for instance represent operations as elements of the symplectic group acting on a list of creation and annihilation operators [Arv+95; TMM21]. However, these methods are mainly geared towards handling the complexity of e.g. Gaussian states in continuous-variable quantum computing, and seem to be of little use for the calculations of interest here. Instead, this chapter describes a "component-wise" way to express the time-dependent ion–photon wavepacket calculations that can easily be evaluated using virtually any symbolic or numerical computer systems, that is, without having to explicitly define symbolic time-dependent annihilation and creation operators and their commutation relations.

We first identify single-photon excitations of a single spatial/polarisation mode with their temporal envelope function (in the notation of eq. (2.41), $B_{n\Pi}^\dagger |0\rangle$ with β). We drop the requirement for normalisation, and consider the Hilbert space of such single-mode excitations to be that of square-integrable functions $L_{\mathbb{C}}^2(\mathbb{R})$ with the natural scalar product $\langle f, g \rangle = \int_{-\infty}^{\infty} \overline{f(t)} g(t) \, \mathrm{d}t$. We write $\mathcal{A}_{n\Pi}$ for this space in its role for excitations for the mode $n\Pi$ (though they are of course all isomorphic). The full Fock space for an arbitrary number of excitations in that mode would be given by $F(\mathcal{A}_{n\Pi}) = \bigoplus_{n=0}^{\infty} S\mathcal{A}_{n\Pi}^{\otimes n} = \mathbb{C} \oplus \mathcal{A}_{n\Pi} \oplus S(\mathcal{A}_{n\Pi} \otimes \mathcal{A}_{n\Pi}) \oplus S(\mathcal{A}_{n\Pi} \otimes \mathcal{A}_{n\Pi} \otimes \mathcal{A}_{n\Pi}) \oplus \cdots$ (or, more precisely, the closure of the same), where S is the symmetrisation operator. The multi-mode Hilbert space would then in turn be given by the tensor product of the single-mode Fock spaces, $\mathcal{H}_P = \bigotimes_{n\Pi} F(\mathcal{A}_{n\Pi})$.

As we consider only up to two photons in the system (one from each ion, with the apparatus preserving the photon number), we can truncate the full Hilbert space to the components with no more than two excitations. For two modes $n \in \{A, B\}$ and polari-

sations $\Pi \in \{H, V\}$, this two-excitation-truncated space $\mathcal{H}_P^{\leq 2}$ would be:

$$\mathcal{H}_P^{\leq 2} := \underbrace{\mathbb{C}}_{=:\mathcal{H}_P^{(0)},\text{ zero photons}} \oplus \underbrace{\mathcal{A}_{AH} \oplus \mathcal{A}_{AV} \oplus \mathcal{A}_{BH} \oplus \mathcal{A}_{BV}}_{=:\mathcal{H}_P^{(1)},\text{ one photon}} \oplus$$

$$\underbrace{(\mathcal{A}_{AH} \otimes \mathcal{A}_{AH}) \oplus (\mathcal{A}_{AH} \otimes \mathcal{A}_{AV}) \oplus (\mathcal{A}_{AH} \otimes \mathcal{A}_{BH}) \oplus (\mathcal{A}_{AH} \otimes \mathcal{A}_{BV}) \oplus (\mathcal{A}_{AV} \otimes \mathcal{A}_{AV}) \oplus}_{}$$
$$\underbrace{(\mathcal{A}_{AV} \otimes \mathcal{A}_{BH}) \oplus (\mathcal{A}_{AV} \otimes \mathcal{A}_{BV}) \oplus (\mathcal{A}_{BH} \otimes \mathcal{A}_{BH}) \oplus (\mathcal{A}_{BH} \otimes \mathcal{A}_{BV}) \oplus (\mathcal{A}_{BV} \otimes \mathcal{A}_{BV})}_{=:\mathcal{H}_P^{(2)},\text{ two photons}}. \quad \text{(A.1)}$$

For notational uniformity throughout this treatment, we have elided the explicit symmetrisation operator for the two-photon subspaces $\mathcal{A}_{AH} \otimes \mathcal{A}_{AH}, \ldots$ (where there are two excitations in the same mode); we will explicitly ensure that the states and measurements remain symmetric[1]. Expansion including additional modes (e.g. modes C and D in fig. 4.21) would proceed much the same way.

The pure ion–photon input state from eq. (4.46) for system $n \in \{A, B\}$ would then lie in $\mathcal{H}_I \otimes (\mathcal{A}_{nH} \oplus \mathcal{A}_{nV}) \subset \mathcal{H}_I \otimes \mathcal{H}_P^{\leq 2}$, where $\mathcal{H}_I := \mathbb{C}^2$ are the relevant ion qubit levels for either of the ions. The support of the joint state of the two ion–photon systems, given by the tensor product $|\psi_{IP}^A\rangle \otimes |\psi_{IP}^B\rangle$, is all of $\mathcal{H}_I^{\otimes 2} \otimes \mathcal{H}_P^{(2)}$. We will simply write $\mathcal{H}$ for the state space of the full system, $\mathcal{H} := \mathcal{H}_I^{\otimes 2} \otimes \mathcal{H}_P^{\leq 2}$.

Equation (A.1) suggests a straightforward "enumerated" representation of a state $|\zeta\rangle \in \mathcal{H}_P^{\leq k}$ as a vector of complex-valued coefficient functions

$$(\zeta_\varnothing, \zeta_{AH}(t), \zeta_{AV}(t), \ldots, \zeta_{AH,AH}(t, t'), \zeta_{AH,AV}(t, t'), \ldots)^T, \quad \text{(A.2)}$$

where the number of time parameters varies with the number of total excitations. Their ordering requires an arbitrary choice; a natural convention is to order the entries lexicographically in terms of their mode index tuples $(n\Pi, n'\Pi', \ldots)$, which are themselves sorted lexicographically, and then to consider the time indices $\{t, t', t'', \ldots\}$ to always apply to those modes from left to right (note that the entries of the vector are the functions themselves, not their values when evaluated at a particular time; the arguments were just listed in eq. (A.2) for illustration). This is just a way to define one conventional order for the commutative (under a canonical isomorphism) tensor product, and an order when writing such entries in a vector; it is not, for instance, any statement about the bosonic nature of the excitations. In general, we will just invoke the permutation symmetries when determining the matrix representation for operators on this space, but not in the vector space structure itself.

Both the "photonic part" of the input density operator $\tau_{i^A j^A}^A \otimes \tau_{i^B j^B}^B \in \mathcal{S}(\mathcal{H}_P^{\leq 2})$ from eq. (4.48) and the unitary describing the action of the beamsplitter network U_P (restricted to the excitation-number-truncated Hilbert space) are linear operators in $\mathrm{End}(\mathcal{H}_P^{\leq 2}) \cong \mathcal{H}_P^{\leq 2} \otimes (\mathcal{H}_P^{\leq 2})^*$. As such, we can consider them as square matrices with each axis arranged

---

[1] As the input states to the beamsplitter network consist of only one-photon terms, this is trivial, and on the output side, two photons in the same mode turn out not to be interesting anyway.

in the same vein as in eq. (A.2). To keep track of the time parameters from the different sides, we will use $\{l, l', l'', \ldots\}$ for the "left" parameters corresponding to creation operators/kets, $\{r, r', r'', \ldots\}$ for the "right" parameters corresponding to annihilation operators/bras. For instance, the element $\xi_{AH,BH;AV}(l, l'; r)$ of a matrix ξ would be the term acting as $\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \xi_{AH,BH;AV}(l, l'; r) \, b_{AH}^{\dagger}(l) b_{BH}^{\dagger}(l') \, |0\rangle\langle 0| \, b_{AV}(r) \, \mathrm{d}r\mathrm{d}l'\mathrm{d}l$. Matrix multiplication is defined in the natural way as a pointwise product of functions where the "inner" time parameters ($\{r, r', \ldots\}$ of the left operand, $\{l, l', \ldots\}$ of the right operand) are matched and integrated over, leaving a function in the "outer" parameters ($\{l, l', \ldots\}$ of the left operand, $\{r, r', \ldots\}$ of the right operand), matching what is anyway required by dimensionality considerations. A column vector such as that from eq. (A.2), matching the usual isomorphic interpretation as a $k \times 1$ matrix, has only "left" time parameters.

A.1   Lifting of two-mode unitary operations

The action of a number of optical elements (beamsplitters, birefringent elements) on a photonic state given in second quantisation can be expressed by substituting the creation operators $\{c^{\dagger}, d^{\dagger}\}$ of the two involved modes according to a 2×2 unitary matrix U,

$$\begin{pmatrix} c^{\dagger} \\ d^{\dagger} \end{pmatrix} \mapsto U^{T} \begin{pmatrix} c^{\dagger} \\ d^{\dagger} \end{pmatrix}, \tag{A.3}$$

where the transpose ensures that U is the same entry by entry as the unitary matrix giving the qubit operation in the one-excitation subspace of the participating two modes, span $\{c^{\dagger} |0\rangle, d^{\dagger} |0\rangle\}$.

In the present context, the transformations act on a larger system with a number of other "bystander" modes, and multiple excitations (here, up to two). We assume the transformations to be stationary and sufficiently wavelength-independent such that they just act on two same-time creation operators $b_{n\Pi}^{\dagger}(t)$ and $b_{n'\Pi'}^{\dagger}(t)$ independently of t. We can then straightforwardly embed, or *lift*, the two-mode unitary matrix into the space of operators on $\mathcal{H}_P^{\leq 2}$. We will denote this lifting process for two modes x and y (specified by spatial mode and polarisation; $x \neq y$) as $\mathcal{V}_{x \leftrightarrow y} : \mathrm{U}(2) \to \mathrm{End}(\mathcal{H}_P^{\leq 2}), U \mapsto \mathcal{V}_{x \leftrightarrow y}(U)$. The explicit form of the lifted matrix $\mathcal{V}_{x \leftrightarrow y}(U)$ is easily derived by inserting a decomposition of the identity on $\mathcal{H}_P^{\leq 2}$ on either side of $\mathcal{V}_{x \leftrightarrow y}(U)$ and evaluating the resulting matrix elements using the canonical commutation relations and action on the vacuum state as usual. The resulting matrix "factors into blocks" in the following sense:

Consider $U \in \mathrm{U}(2)$ to be of the form

$$U = \begin{pmatrix} \alpha & \beta \\ \gamma & \varepsilon \end{pmatrix} \tag{A.4}$$

for coefficients $\alpha, \beta, \gamma, \varepsilon \in \mathbb{C}$, and assume without loss of generality that $x < y$ in lexicographic order (U can just be flipped otherwise). On the subspace of $\mathcal{H}_P^{\leq 2}$ which contains no

excitations in the modes x or y, $\mathcal{V}_{x\leftrightarrow y}(U)$ acts as the identity, or in terms of the amplitude functions, merely "forwards" the time arguments. The zero-photon term is 1, one-photon terms are $\delta(l-r)$, two-photon terms $\delta(l-r)\delta(l'-r')$, etc. For instance, if we take the truncated four-mode space from eq. (A.1) and $x = AH$, $y = BH$, then these non-participating terms are $\mathbb{C}$, $\mathcal{A}_{AV}$, $\mathcal{A}_{BV}$, $\mathcal{A}_{AV} \otimes \mathcal{A}_{AV}$, $\mathcal{A}_{AV} \otimes \mathcal{A}_{BV}$, and $\mathcal{A}_{BV} \otimes \mathcal{A}_{BV}$.

For subspaces of $\mathcal{H}_P^{\leq 2}$ containing one excitation in x or y, $\mathcal{V}_{x\leftrightarrow y}(U)$ acts in the straightforward fashion by forming the linear combination between the respective terms with x and those in y according to the elements of U. For instance, the matrix elements of $\mathcal{V}_{x\leftrightarrow y}(U)$ restricted to $\mathcal{A}_x = \mathcal{A}_{AH}$ and $\mathcal{A}_y = \mathcal{A}_{BH}$ are just

$$\begin{pmatrix} \alpha\,\delta(l-r) & \beta\,\delta(l-r) \\ \gamma\,\delta(l-r) & \varepsilon\,\delta(l-r) \end{pmatrix}. \tag{A.5}$$

Note that these are not the only terms with matrix elements of this form. Generally, considering all the terms containing one of $\mathcal{A}_{AH}$ or $\mathcal{A}_{BH}$ we can collect them as $\mathcal{A}_{AH} \otimes (\mathbb{C} \oplus \mathcal{A}_{AV} \oplus \mathcal{A}_{BV})$ and $\mathcal{A}_{BH} \otimes (\mathbb{C} \oplus \mathcal{A}_{AV} \oplus \mathcal{A}_{BV})$, where $\mathcal{V}_{x\leftrightarrow y}$ then acts on each of the blocks transversally like in eq. (A.5), carrying through any extra time parameters. In pulling the tensor products to the front like this, we have identified $\mathcal{A}_{AH} \cong \mathcal{A}_{AH} \otimes \mathbb{C}$. We also have made use of the commutative property of the tensor product, as e.g. distributing tensor product in $\mathcal{A}_{BH}$ yields $\mathcal{A}_{BH} \otimes \mathcal{A}_{AH}$, contrary to the "canonical" lexicographic order. Such order swaps will result in a corresponding permutation of time parameters, giving e.g. $\delta(l-r')\delta(l'-r)$ factors.

Finally, we have subspaces containing two excitations of the modes x and y. We chose the cut-off in $\mathcal{H}_P^{\leq k}$ at $k = 2$, so these are the highest-order terms, which also only consist of $\mathcal{A}_{AH} \otimes \mathcal{A}_{AH}$, $\mathcal{A}_{AH} \otimes \mathcal{A}_{BH}$ and $\mathcal{A}_{BH} \otimes \mathcal{A}_{BH}$. Restricted to those three elements, the action of $\mathcal{V}_{x\leftrightarrow y}(U)$, resulting from the canonical commutation relations, is (writing $\delta_{\rightrightarrows} := \delta(l-r)\delta(l'-r')$, $\delta_{\rtimes} := \delta(l-r')\delta(l'-r)$):

$$\begin{pmatrix} \alpha^2\,\delta_{\rightrightarrows} & \alpha\beta\,\delta_{\rightrightarrows} & \beta^2\,\delta_{\rightrightarrows} \\ \alpha\gamma(\delta_{\rightrightarrows} + \delta_{\rtimes}) & \alpha\varepsilon\,\delta_{\rightrightarrows} + \beta\gamma\,\delta_{\rtimes} & \beta\varepsilon(\delta_{\rightrightarrows} + \delta_{\rtimes}) \\ \gamma^2\,\delta_{\rightrightarrows} & \gamma\varepsilon\,\delta_{\rightrightarrows} & \varepsilon^2\,\delta_{\rightrightarrows} \end{pmatrix}. \tag{A.6}$$

For higher excitation number cutoffs, this would also apply across blocks with other bystander modes in the same fashion as illustrated for the one-excitation case here.

To complete the explicit example, the full matrix of $\mathcal{V}_{AH\leftrightarrow BH}(U)$ in the truncated four-mode, two-excitation space from eq. (A.1) obtained by following the above prescriptions

is[2]:

$$
\left[
\begin{array}{cccc:cccccccccccc}
1 \\
\alpha\delta(l-r) & 0 & \beta\delta(l-r) & 0 \\
0 & \delta(l-r) & 0 & 0 \\
\gamma\delta(l-r) & 0 & \varepsilon\delta(l-r) & 0 \\
0 & 0 & 0 & \delta(l-r) \\
\hdashline
 & & & & \alpha^2\delta_\rightrightarrows & 0 & \alpha\beta\delta_\rightrightarrows & 0 & 0 & 0 & 0 & \beta^2\delta_\rightrightarrows & 0 & 0 \\
 & & & & 0 & \alpha\delta_\rightrightarrows & 0 & 0 & 0 & \beta\delta_\searrow & 0 & 0 & 0 & 0 \\
 & & & & \alpha\gamma(\delta_\rightrightarrows+\delta_\searrow) & 0 & \alpha\varepsilon\delta_\rightrightarrows+\beta\gamma\delta_\searrow & 0 & 0 & 0 & 0 & \beta\varepsilon(\delta_\rightrightarrows+\delta_\searrow) & 0 & 0 \\
 & & & & 0 & 0 & 0 & \alpha\delta_\rightrightarrows & 0 & 0 & 0 & 0 & \beta\delta_\rightrightarrows & 0 \\
 & & & & 0 & 0 & 0 & 0 & \delta_\rightrightarrows & 0 & 0 & 0 & 0 & 0 \\
 & & & & 0 & \gamma\delta_\searrow & 0 & 0 & 0 & \varepsilon\delta_\rightrightarrows & 0 & 0 & 0 & 0 \\
 & & & & 0 & 0 & 0 & 0 & 0 & 0 & \delta_\rightrightarrows & 0 & 0 & 0 \\
 & & & & \gamma^2\delta_\rightrightarrows & 0 & \gamma\varepsilon\delta_\rightrightarrows & 0 & 0 & 0 & 0 & \varepsilon^2\delta_\rightrightarrows & 0 & 0 \\
 & & & & 0 & 0 & 0 & \gamma\delta_\rightrightarrows & 0 & 0 & 0 & 0 & \varepsilon\delta_\rightrightarrows & 0 \\
 & & & & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & \delta_\rightrightarrows
\end{array}
\right].
$$

Because the action of U preserves the number of excitations, the off-diagonal blocks that would mix states of different excitation numbers are all zero.

A.2   Ideal ion–photon input states

For two ideal ion–photon states according to eq. (4.46), which we assume to arrive aligned in time (for simplicity, $\Delta t^A = \Delta t^B = 0$), the relevant non-zero coefficients for the truncated Fock space representation ξ of the photonic part $\tau^A_{i^A j^A} \otimes \tau^B_{i^B j^B}$ are all in the two-excitation subspace (discarding the non-coupled parts as we will post-select on the presence of two photons, as explained in §4.4.1), and are given by

$$
\begin{aligned}
\xi_{A o^A, B o^B; A p^A, B p^B}(l, l'; r, r') &= \rho^A_{i^A j^A o^A p^A}(l, r)\, \rho^B_{i^B j^B o^B p^B}(l', r') \\
&= \frac{1}{4}\, \delta_{\pi(i^A) o^A}\, \delta_{\pi(j^A) p^A}\, \delta_{\pi(i^B) o^B}\, \delta_{\pi(j^B) p^B}\, \phi_{i^A}(l)\, \phi_{i^B}(l')\, \overline{\phi_{j^A}(r)}\, \overline{\phi_{j^B}(r')},
\end{aligned}
\tag{A.7}
$$

where the first equality is the generic expansion of the tensor product, and we have inserted the ideal states (4.46) for the second.

A.3   Timestamped photon detection operators

The last ingredient necessary to implement the post-selected state computation from eq. (4.48) is the representation of the measurement operators in the truncated Fock space formalism. This is straightforward:

---

[2]Recall that the elements of this matrix (or its equivalent for other sets of modes) are from the ring(s) of amplitude functions, such that the element multiplication operation when taking a matrix–matrix product consists of integrating over the "inner" time parameters.

Detection of two photons in modes x, y (for $x \neq y$) at times t and t' is given by the POVM element $(\eta_{\text{det}})^2 b_x^\dagger(t) b_y^\dagger(t') |0\rangle\langle 0| b_y(t') b_x(t)$. Its matrix representation is simply the matrix that has one non-zero element on the diagonal,

$$\xi_{x,y;x,y}(l, l'; r, r') = (\eta_{\text{det}})^2 \delta(l - t)\, \delta(l' - t')\, \delta(r - t)\, \delta(r' - t'), \qquad (A.8)$$

and is zero elsewhere. For the detection of two photons in the same mode x, the one non-zero element acquires a second term with t and t' swapped as a result of the indistinguishability of the excitations:

$$
\begin{aligned}
\xi_{x,x;x,x}(l, l'; r, r') = (\eta_{\text{det}})^2 \Big(& \delta(l - t)\, \delta(l' - t')\, \delta(r - t)\, \delta(r' - t') + \\
& \delta(l - t')\, \delta(l' - t)\, \delta(r - t')\, \delta(r' - t) \Big).
\end{aligned}
\qquad (A.9)
$$

In the general case, the trace operator in eq. (4.48) is equivalent to a summation of the diagonal entries of the truncated Fock space matrices, along with an integration over the respective pairs of time parameters $l \leftrightarrow r$, $l' \leftrightarrow r'$, etc. However, note that if F is a POVM operator of the form (A.8), multiplying by F and taking the trace is equivalent to just evaluating the $(x, y; x, y)$ element for $l = r = t$, $l' = r' = t'$ (plus the swapped assignment for the same mode according to eq. (A.9)). Similarly, for all the other integrations from the matrix-matrix products, at least one of the functions is a linear combination of delta functions. Thus, these only consist of a propagation of coefficients associated with function evaluation at given time parameters as well, and no integrations actually need to be carried out for the evaluation of eq. (4.48), as expected from a theory where interactions are instantaneous in time.

Evidently, systems of larger number of modes and operations will still preferentially be manipulated using automated means. However, the formalism presented here provides a straightforward way to implement the calculations in any standard numerical or symbolic programming environment. The evaluation of any integrals is only necessary in situations where they are intrinsic to the problem, e.g. when averaging over a range of possible photon arrival times.

Glossary of notation, symbols and abbreviations

Mathematical notation

$\overline{\cdot}$	complex conjugation
$\cdot^T$	transpose of a vector or matrix
$\cdot^\dagger$	Hermitian conjugate (conjugate transpose)
$\lvert\cdot\rvert$	absolute value; 2-norm of a vector in $\mathbb{R}^n$ or $\mathbb{C}^n$; cardinality of a set
$\lVert\cdot\rVert$	general 2-norm of a vector, e.g. the L^2-norm of a complex function
$\lVert\cdot\rVert_1$	1-norm (trace norm) of a linear operator ($\lVert A\rVert_1 = \mathrm{tr}\ \sqrt{A^\dagger A}$)
$\cdot$	scalar product of vectors; multiplication of numbers
$\times$	cross product of vectors; multiplication of numbers
$\hat{\cdot}$	Euclidean unit vector (linear operators are not adorned thusly)
$\otimes$	tensor product
$\oplus$	direct sum; addition modulo 2 of bits and bitstrings (xor)
$\lvert\cdot\rangle, \langle\cdot\rvert$	quantum state and associated 1-form in Dirac notation
$\langle\cdot\rangle$	average or expectation value of the given quantity
$\langle\cdot,\cdot\rangle$	inner product of two complex functions
$[\cdot,\cdot]$	commutator of two operators ($[a,b] = ab - ba$)
$\{\cdot,\cdot\}$	a two-element set (does not appear here in reference to an anticommutator)
$\mathbb{C}$	complex numbers
$\mathbb{R}$	real numbers
$\mathbb{Z}$	integers
$\mathbb{N}$	positive integers
$\mathbb{1}$	identity operator
$\arg$	argument of a complex number ($z = \lvert z\rvert e^{i\arg z}$)
Atan	two-argument arctangent function ($\arg z = \mathrm{Atan}(\mathrm{im}(z), \mathrm{re}(z))$)
e	Euler's number, exponential function
i	imaginary unit
$\mathrm{re}(\cdot)$	real part of a complex number
$\mathrm{im}(\cdot)$	imaginary part of a complex number
J_k	Bessel function of the first kind

$O(\cdot)$	"big O" Landau asymptotic notation
$\mathcal{S}(\mathcal{H})$	bounded operators on a Hilbert space $\mathcal{H}$
$SO(n)$	special orthogonal group in n dimensions
span	linear hull of a set of vectors
tr	trace of linear operators
tr_B	partial trace eliminating subsystem B
$U(n)$	unitary group in n dimensions
$\delta(t), \delta(\boldsymbol{v})$	Dirac delta distribution (normalised to unity)
δ_{mn}	Kronecker delta (1 if $m = n$, 0 otherwise)
$\Theta(t)$	Heaviside unit step function (1 if $t \geq 0$, 0 otherwise)
$\boldsymbol{a}, \boldsymbol{b}, \dots$	vectors; messages from a two- or three-character alphabet (in chapter 7)
$\hat{\boldsymbol{r}}$	normalised direction of r ($\hat{\boldsymbol{r}} = \boldsymbol{r}/|r|$)
$\nabla\Phi$	gradient of a scalar field Φ
$\boldsymbol{H}_\Phi$	Hessian matrix of a scalar field Φ
$\nabla\boldsymbol{v}$	gradient of a vector field $\boldsymbol{v}$
$\nabla \cdot \boldsymbol{v}$	divergence of a vector field $\boldsymbol{v}$
c. c.	complex conjugate
h. c.	Hermitian conjugate

$$\approx$$

Commonly used symbols

$a_{\boldsymbol{k}\Pi}^\dagger$	plane-wave creation operator for angular wavevector $\boldsymbol{k}$ and polarisation Π
$b_{n\Pi}^\dagger(t)$	time-domain creation operator for beam mode n and polarisation Π (see §2.3.1)
$\boldsymbol{d}_{ie}$	atomic dipole for the transition $|i\rangle \leftrightarrow |e\rangle$
e	elementary charge ($e := 1.602\,176\,634 \times 10^{-19}$ C)
$\hat{\boldsymbol{e}}_0, \hat{\boldsymbol{e}}_1, \hat{\boldsymbol{e}}_{-1}$	spherical basis vectors
$\mathcal{E}$	Bell-state error/infidelity ($1 - \mathcal{F}$)
$\mathcal{F}$	fully entangled fraction ("Bell-state fidelity", see §2.1.3)
g_S	g-factor of the electron ($g_S \approx 2.002319$)
h	trap RF amplitude modulation index (in chapter 5); binary entropy function (in chapter 7)
$\tilde{h}$	rescaled trap RF amplitude modulation index (eq. (5.4))
$\hbar$	reduced Planck constant ($2\pi\hbar := 6.626\,070\,15 \times 10^{-34}$ J/Hz)
$\mathcal{H}$	Hilbert space of quantum states
H	Hamiltonian describing a quantum system
$H(\rho_A)$	von-Neumann entropy of the state ρ_A

$H(A)_\rho$	alternative notation for $H(\rho_A)$
$H(A\vert B)_\rho$	conditional entropy of system A given B for state ρ
$\boldsymbol{k}$	angular wavevector
k	angular wavenumber
k_B	Boltzmann constant ($\mathrm{k}_B \coloneqq 1.380\,649 \times 10^{-23}$ J/K)
$\mathcal{L}$	likelihood function
$\overline{n}$	mean excitation number of a harmonic oscillator mode
$\cdot^s$	superscript denoting system/mode in chapter 4 ($s \in \{A, B\}$ or $\{A, B, C, D\}$)
$\mathcal{V}_{a\leftrightarrow b}(U)$	unitary U acting on modes a and b (see appendix A)
$\hat{\boldsymbol{x}}, \hat{\boldsymbol{y}}, \hat{\boldsymbol{z}}$	Cartesian basis vectors
Z	ion charge, in units of e
ε_0	permittivity of free space
μ_0	permeability of free space
μ_B	Bohr magneton
Π	transversal polarisation index
ρ_{AB}	density matrix for a composite system with subsystems A and B
$\rho_{\wedge X}$	sub-normalised density matrix conditioned on event X
$\rho_{\vert X}$	normalised density matrix conditioned on event X, $\rho_{\vert X} = \rho_{\wedge X}/\operatorname{tr} \rho_{\wedge X}$
Φ_{dc}	static potential in a Paul trap
Φ_{rf}	amplitude of instantaneous oscillating potential in a Paul trap
$\Phi_{\overline{\mathrm{rf}}}$	pseudopotential describing the RF-induced confinement in a Paul trap
$\phi_i(t)$	temporal envelope of photon wavepacket for decay to $\vert i\rangle$ (see eq. (4.3))
ω	winning probability of the CHSH game (in chapter 7)
Ω_{rf}	angular frequency of the oscillating field in a Paul trap
ω_0	angular frequency of a secular mode of motion
ω_{am}	angular frequency at which the RF trap drive is modulated (in chapter 5)

$\approx$

Acronyms and abbreviations

AC	alternating current (such as that of the electrical power grid)
AOM	acousto-optic modulator
APD	avalanche photodiode
CHSH	Clauser–Horne–Shimony–Holt (authors of ref. [Cla+69])
CNOT	the conditional NOT quantum gate ($\mathbb{1} \oplus X$)
CPTP	completely positive and trace-preserving
CZ	the conditional Z quantum gate ($\mathbb{1} \oplus Z$)

DAC	digital-to-analogue converter
DC	direct current (colloquially used to refer to static, non-oscillatory potentials)
DDS	direct-digital synthesis (a method for radio-frequency generation)
DIQKD	device-independent quantum key distribution
FPGA	field-programmable gate array
HF	higher-frequency (… radial mode of motion of a single ion)
i.i.d	independent and identically distributed
IP	in-plane (parallel to the surface of a planar ion trap)
LF	lower-frequency (… radial mode of motion of a single ion)
LOCC	local operations and classical control
MEMS	micro-electromechanical system
MLE	maximum-likelihood estimation
NA	numerical aperture
NV	nitrogen vacancy (a type of optically active point defect in diamond)
OOP	out-of-plane (orthogonal to the surface of a planar ion trap)
PBS	polarising beamsplitter
PC	personal computer
PMT	photomultiplier tube
POVM	positive operator-valued measure
QBER	quantum bit-error rate
QKD	quantum key distribution
QRNG	quantum random number generator
RF	radio frequency
TTL	transistor–transistor logic (conventionally used to refer to a binary high/low voltage signal)

≈

Bibliography

[Ací+07]   A. Acín, N. Brunner, N. Gisin, S. Massar, S. Pironio, and V. Scarani, "Device-Independent Security of Quantum Cryptography against Collective Attacks", Physical Review Letters **98**, 230501 (2007).

[AGM06]   A. Acín, N. Gisin, and L. Masanes, "From Bell's Theorem to Secure Quantum Key Distribution", Physical Review Letters **97**, 120405 (2006).

[AGR81]   A. Aspect, P. Grangier, and G. Roger, "Experimental Tests of Realistic Local Theories via Bell's Theorem", Physical Review Letters **47**, 460–463 (1981).

[Agu+16]   E. A. Aguilar, R. Ramanathan, J. Kofler, and M. Pawłowski, "Completely device-independent quantum key distribution", Physical Review A **94**, 022305 (2016).

[Ake+10]   N. Akerman, S. Kotler, Y. Glickman, Y. Dallal, A. Keselman, and R. Ozeri, "Single-ion nonlinear mechanical oscillator", Physical Review A **82**, 061402 (2010).

[Ald+16]   I. Alda, J. Berthelot, R. A. Rica, and R. Quidant, "Trapping and manipulation of individual nanoparticles in a planar Paul trap", Applied Physics Letters **109**, 163105 (2016).

[Alj+13]   S. A. Aljunid, G. Maslennikov, Y. Wang, H. L. Dao, V. Scarani, and C. Kurtsiefer, "Excitation of a Single Atom with Exponentially Rising Light Pulses", Physical Review Letters **111**, 103001 (2013).

[All+10]   D. T. C. Allcock, J. A. Sherman, D. N. Stacey, A. H. Burrell, M. J. Curtis, G. Imreh, N. M. Linke, D. J. Szwer, S. C. Webster, A. M. Steane, and D. M. Lucas, "Implementation of a symmetric surface-electrode ion trap with field compensation using a modulated Raman effect", New Journal of Physics **12**, 053026 (2010).

[AM84]   J. R. Ackerhalt and P. W. Milonni, "Interaction Hamiltonian of quantum optics", Journal of the Optical Society of America B **1**, 116 (1984).

[And+18]   D. L. Andrews, G. A. Jones, A. Salam, and R. G. Woolley, "Perspective: Quantum Hamiltonians for optical interactions", The Journal of Chemical Physics **148**, 040901 (2018).

[And+20]   D. L. Andrews, D. S. Bradshaw, K. A. Forbes, and A. Salam, "Quantum electrodynamics in modern optics and photonics: tutorial", Journal of the Optical Society of America B **37**, 1153 (2020).

[Ara+18]   G. Araneda, S. Walser, Y. Colombe, D. B. Higginbottom, J. Volz, R. Blatt, and A. Rauschenbeutel, "Wavelength-scale errors in optical localization due to spin–orbit coupling of light", Nature Physics **15**, 17–21 (2018).

[Ara+20]   G. Araneda, G. Cerchiari, D. B. Higginbottom, P. C. Holz, K. Lakhmanskiy, P. Obšil, Y. Colombe, and R. Blatt, "The Panopticon device: An integrated Paultrap–hemispherical mirror system for quantum optics", Review of Scientific Instruments **91**, 113201 (2020).

[Arn+18]   R. Arnon-Friedman, F. Dupuis, O. Fawzi, R. Renner, and T. Vidick, "Practical device-independent quantum cryptography via entropy accumulation", Nature Communications **9**, 459 (2018).

[Arn+19]   K. J. Arnold, R. Kaewuan, T. R. Tan, and M. D. Barrett, "Oscillating quadrupole effects in high-precision metrology", Physical Review A **99**, 022515 (2019).

[Arv+95]   Arvind, B. Dutta, N. Mukunda, and R. Simon, "The real symplectic groups in quantum mechanics and optics", Pramana **45**, 471–497 (1995).

[ARV19]   R. Arnon-Friedman, R. Renner, and T. Vidick, "Simple and Tight Device-Independent Security Proofs", SIAM Journal on Computing **48**, 181–225 (2019).

[Bab84]   M. Babiker, "Theory of momentum changes in atomic translation due to irradiation with resonant light", Journal of Physics B: Atomic and Molecular Physics **17**, 4877–4883 (1984).

[Bad+00]   P. Badziąg, M. Horodecki, P. Horodecki, and R. Horodecki, "Local environment can enhance fidelity of quantum teleportation", Physical Review A **62**, 012311 (2000).

[Bal+16]   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).

[Bal+18]   T. G. Ballance, J. F. Goodwin, B. Nichol, L. J. Stephenson, C. J. Ballance, and D. M. Lucas, "A short response time atomic source for trapped ion experiments", Review of Scientific Instruments **89**, 053102 (2018).

[Ban21]   J.-D. Bancal, *UVMAC source code*, available at `https : / / github . com / jdbancal/uvmac`, 2021.

[BB84]   C. H. Bennett and G. Brassard, "Quantum cryptography: public key distri-
            bution and coin tossing", in IEEE International Conference on Computers,
            Systems and Signal Processing, Vol. 175 (1984), p. 8.

[BBL93]   C. Baxter, M. Babiker, and R. Loudon, "Canonical approach to photon pres-
            sure", Physical Review A **47**, 1278–1287 (1993).

[BCK13]   J. Barrett, R. Colbeck, and A. Kent, "Memory Attacks on Device-Independent
            Quantum Cryptography", Physical Review Letters **110**, 010503 (2013).

[Bel64]   J. S. Bell, "On the Einstein Podolsky Rosen paradox", Physics Physique Fizika
            1, 195–200 (1964).

[Ben+07]   J. Benhelm, G. Kirchmair, U. Rapol, T. Körber, C. F. Roos, and R. Blatt, "Mea-
            surement of the hyperfine structure of the $S_{1/2} - D_{5/2}$ transition in $^{43}Ca^+$",
            Physical Review A **75**, 032506 (2007).

[Ben+96]   C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters, "Mixed-
            state entanglement and quantum error correction", Physical Review A **54**,
            3824 (1996).

[Ber+13]   H. Bernien, B. Hensen, W. Pfaff, G. Koolstra, M. S. Blok, L. Robledo, T. H.
            Taminiau, M. Markham, D. J. Twitchen, L. Childress, and R. Hanson, "Her-
            alded entanglement between solid-state qubits separated by three metres",
            Nature **497**, 86–90 (2013).

[Ber+98]   D. J. Berkeland, J. D. Miller, J. C. Bergquist, W. M. Itano, and D. J. Wineland,
            "Minimization of ion micromotion in a Paul trap", Journal of Applied Physics
            83, 5025–5033 (1998).

[Ber02]   D. J. Berkeland, "Linear Paul trap for strontium ions", Review of Scientific
            Instruments **73**, 2856–2860 (2002).

[BHK05]   J. Barrett, L. Hardy, and A. Kent, "No Signaling and Quantum Key Distribu-
            tion", Physical Review Letters **95**, 010503 (2005).

[BKF09]   A. Broadbent, E. Kashefi, and J. Fitzsimons, "Universal blind quantum com-
            putation", Proceedings of the 50th Annual IEEE Symposium on Foundations
            of Computer Science **6154 LNCS**, 517–526 (2009).

[BKF10]   A. Broadbent, E. Kashefi, and J. F. Fitzsimons, "Measurement-Based and Uni-
            versal Blind Quantum Computation", in *Formal Methods for Quantitative As-
            pects of Programming Languages*, edited by A. Aldini, M. Bernardo, A. Di
            Pierro, and H. Wiklicky, LNCS 6154 (Springer, Berlin, Heidelberg, 2010).

[Bli+04]   B. B. Blinov, D. L. Moehring, L.-M. Duan, and C. Monroe, "Observation of
            entanglement between a single trapped atom and a single photon", Nature
            428, 153–157 (2004).

[Blu10] R. Blume-Kohout, "Optimal, reliable estimation of quantum states", New Journal of Physics **12**, 043034 (2010).

[Blu12] R. Blume-Kohout, *Robust error bars for quantum tomography*, (Feb. 23, 2012) preprint.

[BP01] S. Brouard and J. Plata, "Heating of a trapped ion by random fields: The influence of the micromotion", Physical Review A **63**, 043402 (2001).

[BPT74] M. Babiker, E. A. Power, and T. Thirunamachandran, "On a generalization of the Power–Zienau–Woolley transformation in quantum electrodynamics and atomic field equations", Proc. R. Soc. Lond. A **338**, 245–449 (1974).

[Bre+15] B. Brecht, D. V. Reddy, C. Silberhorn, and M. G. Raymer, "Photon Temporal Modes: A Complete Framework for Quantum Information Science", Physical Review X **5**, 041017 (2015).

[Bre+19] S. M. Brewer, J.-S. Chen, A. M. Hankin, E. R. Clements, C. W. Chou, D. J. Wineland, D. B. Hume, and D. R. Leibrandt, "^{27}Al$^+$ Quantum-Logic Clock with a Systematic Uncertainty below 10^{-18}", Physical Review Letters **123**, 033201 (2019).

[Bro20] R. G. Brown, *Dieharder: A Random Number Test Suite*, version 3.31.1, available at `https://webhome.phy.duke.edu/~rgb/General/dieharder.php`, 2020.

[Bru+14] N. Brunner, D. Cavalcanti, S. Pironio, V. Scarani, and S. Wehner, "Bell nonlocality", Reviews of Modern Physics **86**, 419–478 (2014).

[BS68] D. M. Brink and G. R. Satchler, *Angular momentum*, 2nd ed. (Oxford University Press, 1968).

[Bur+19] S. C. Burd, R. Srinivas, J. J. Bollinger, A. C. Wilson, D. J. Wineland, D. Leibfried, D. H. Slichter, and D. T. C. Allcock, "Quantum amplification of mechanical oscillator motion", Science **364**, 1163–1165 (2019).

[Cab+99] C. Cabrillo, J. I. Cirac, P. Garcia-Fernandez, and P. Zoller, "Creation of entangled states of distant atoms by interference", Physical Review A **59**, 10 (1999).

[CB08] E. T. Campbell and S. C. Benjamin, "Measurement-based entanglement under conditions of extreme photon loss", Physical Review Letters **101**, 16–19 (2008).

[CDG97] C. Cohen-Tannoudji, J. Dupont-Roc, and G. Grynberg, *Photons and atoms: introduction to quantum electrodynamics* (Wiley-VCH, Weinheim, 1997), 468 pp.

[Chr+13]   B. G. Christensen, K. T. McCusker, J. B. Altepeter, B. Calkins, T. Gerrits, A. E. Lita, A. Miller, L. K. Shalm, Y. Zhang, S. W. Nam, N. Brunner, C. C. W. Lim, N. Gisin, and P. G. Kwiat, "Detection-Loophole-Free Test of Quantum Nonlocality, and Applications", Physical Review Letters **111**, 130406 (2013).

[Chr+20]   J. E. Christensen, D. Hucul, W. C. Campbell, and E. R. Hudson, "High-fidelity manipulation of a qubit enabled by a manufactured nucleus", npj Quantum Information **6**, 35 (2020).

[Chu+13]   B. L. Chuah, N. C. Lewty, R. Cazan, and M. D. Barrett, "Detection of ion micromotion in a linear Paul trap with a high finesse cavity", Optics Express **21**, 10632 (2013).

[Chw09]   M. Chwalla, "Precision spectroscopy with ^{40}Ca$^+$ ions in a Paul trap", PhD thesis (University of Innsbruck, Innsbruck, 2009).

[Cir+97]   J. I. Cirac, P. Zoller, H. J. Kimble, and H. Mabuchi, "Quantum State Transfer and Entanglement Distribution among Distant Nodes in a Quantum Network", Physical Review Letters **78**, 3221–3224 (1997).

[Cir80]   B. S. Cirel'son, "Quantum generalizations of Bell's inequality", Letters in Mathematical Physics **4**, 93–100 (1980).

[CK18]   B. Coecke and A. Kissinger, "Picturing quantum processes", in International Conference on Theory and Application of Diagrams (2018), pp. 28–31.

[CKW00]   V. Coffman, J. Kundu, and W. K. Wootters, "Distributed entanglement", Physical Review A **61**, 052306 (2000).

[Cla+21]   C. R. Clark, H. N. Tinkey, B. C. Sawyer, A. M. Meier, K. A. Burkhardt, C. M. Seck, C. M. Shappert, N. D. Guise, C. E. Volin, S. D. Fallek, H. T. Hayden, W. G. Rellergert, and K. R. Brown, "High-Fidelity Bell-State Preparation with ^{40}Ca$^+$ Optical Qubits", Physical Review Letters **127**, 130505 (2021).

[Cla+69]   J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt, "Proposed Experiment to Test Local Hidden-Variable Theories", Physical Review Letters **23**, 880–884 (1969).

[Coe10]   B. Coecke, "Quantum picturalism", Contemporary Physics **51**, 59 (2010).

[CR12]   M. Christandl and R. Renner, "Reliable quantum state tomography", Physical Review Letters **109**, 1–6 (2012).

[CW86]   S. Collocott and G. White, "Thermal expansion and heat capacity of some stainless steels and FeNi alloys", Cryogenics **26**, 402–405 (1986).

[Dav74]   G. Davies, "Vibronic spectra in diamond", Journal of Physics C: Solid State Physics **7**, 3797–3809 (1974).

[dCle+16]   L. E. de Clercq, H.-Y. Lo, M. Marinelli, D. Nadlinger, R. Oswald, V. Negnevitsky, D. Kienzler, B. Keitch, and J. P. Home, "Parallel Transport Quantum Logic Gates with Trapped Ions", Physical Review Letters **116**, 080502 (2016).

[De+12]   A. De, C. Portmann, T. Vidick, and R. Renner, "Trevisan's Extractor in the Presence of Quantum Side Information", SIAM Journal on Computing **41**, 915–940 (2012).

[Del+16]   A. Delteil, Z. Sun, W.-b. Gao, E. Togan, S. Faelt, and A. Imamoğlu, "Generation of heralded entanglement between distant hole spins", Nature Physics **12**, 218–223 (2016).

[Del+20]   T. Delord, P. Huillery, L. Nicolas, and G. Hétet, "Spin-cooling of the motion of a trapped diamond", Nature **580**, 56–59 (2020).

[Deu85]   D. Deutsch, "Quantum Theory, the Church-Turing Principle and the Universal Quantum Computer", Proceedings of the Royal Society A **400**, 97–117 (1985).

[DFR20]   F. Dupuis, O. Fawzi, and R. Renner, "Entropy Accumulation", Communications in Mathematical Physics **379**, 867–913 (2020).

[DH76]   W. Diffie and M. E. Hellman, "New Directions in Cryptography", IEEE Transactions on Information Theory **22**, 644–654 (1976).

[DK07]   W. Dai and T. Krovetz, "VHASH Security", Cryptology ePrint Archive (2007).

[DM03]   J. P. Dowling and G. J. Milburn, "Quantum technology: the second quantum revolution", Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences **361**, edited by A. G. J. MacFarlane, 1655–1674 (2003).

[DM10]   L.-M. Duan and C. Monroe, "*Colloquium* : Quantum networks with trapped ions", Reviews of Modern Physics **82**, 1209–1224 (2010).

[DRC17]   C. L. Degen, F. Reinhard, and P. Cappellaro, "Quantum sensing", Reviews of Modern Physics **89**, 035002 (2017).

[Dre+83]   R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward, "Laser phase and frequency stabilization using an optical resonator", Applied Physics B Photophysics and Laser Chemistry **31**, 97–105 (1983).

[Dua+01]   L. M. Duan, M. D. Lukin, J. I. Cirac, and P. Zoller, "Long-distance quantum communication with atomic ensembles and linear optics", Nature **414**, 413–418 (2001).

[Eke18]   A. Ekert, "The power of independence", Nature Physics **14**, 114–115 (2018).

[Eke91]   A. K. Ekert, "Quantum cryptography based on Bell's theorem", Physical Review Letters **67**, 661–663 (1991).

[Elt13]   A. M. Eltony, "Sensitive, 3D micromotion compensation in a surface-electrode ion trap", MSc Thesis (University of British Columbia, 2013).

[End00]   J. Enderlein, "Theoretical study of detection of a dipole emitter through an objective with high numerical aperture", Optics Letters **25**, 634 (2000).

[ER14]   A. Ekert and R. Renner, "The ultimate physical limits of privacy", Nature **507**, 443–447 (2014).

[Esm+21]   I. Esmaeil Zadeh, J. Chang, J. W. N. Los, S. Gyger, A. W. Elshaari, S. Steinhauer, S. N. Dorenbos, and V. Zwiller, "Superconducting nanowire single-photon detectors: A perspective on evolution, state-of-the-art, future developments, and applications", Applied Physics Letters **118**, 190502 (2021).

[FB18]   C. Ferrie and R. Blume-Kohout, *Maximum likelihood quantum state tomography is inadmissible*, (2018) preprint.

[Fed+05]   M. V. Fedorov, M. A. Efremov, A. E. Kazakov, K. W. Chan, C. K. Law, and J. H. Eberly, "Spontaneous emission of a photon: Wave-packet structures and atom-photon entanglement", Physical Review A **72**, 032110 (2005).

[Fey82]   R. P. Feynman, "Simulating physics with computers", International Journal of Theoretical Physics **21**, 467–488 (1982).

[Fis+14]   M. Fischer, M. Bader, R. Maiwald, A. Golla, M. Sondermann, and G. Leuchs, "Efficient saturation of an ion in free space", Applied Physics B **117**, 797–801 (2014).

[FLM19]   N. Funai, J. Louko, and E. Martín-Martínez, "$\hat{p}.\hat{A}$ vs $\hat{x}.\hat{E}$: Gauge invariance in quantum optics and quantum field theory", Physical Review D **99**, 065014 (2019).

[FR16]   P. Faist and R. Renner, "Practical and Reliable Error Bars in Quantum Tomography", Physical Review Letters **117**, 010404 (2016).

[FT20]   C. Fabre and N. Treps, "Modes and states in quantum optics", Reviews of Modern Physics **92**, 035005 (2020).

[Gae+16]   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 Be9+ Ion Qubits", Physical Review Letters **117**, 1–5 (2016).

[Gan+18]   H. C. J. Gan, G. Maslennikov, K.-W. Tseng, T. R. Tan, R. Kaewuam, K. J. Arnold, D. Matsukevich, and M. D. Barrett, "Oscillating-magnetic-field effects in high-precision metrology", Physical Review A **98**, 032514 (2018).

[GAN14]   I. M. Georgescu, S. Ashhab, and F. Nori, "Quantum simulation", Reviews of Modern Physics **86**, 153–185 (2014).

[Gen05]   M. Genovese, "Research on hidden variable theories: A review of recent progresses", Physics Reports **413**, 319–396 (2005).

[Ger+11]   I. Gerhardt, Q. Liu, A. Lamas-Linares, J. Skaar, C. Kurtsiefer, and V. Makarov, "Full-field implementation of a perfect eavesdropper on a quantum cryptography system", Nature Communications **2**, 349 (2011).

[Gha+16]   M. Ghadimi, V. Blūms, B. G. Norton, P. M. Fisher, S. C. Connell, J. M. Amini, C. Volin, H. Hayden, C. S. Pai, D. Kielpinski, M. Lobino, and E. W. Streed, "Scalable ion-photon quantum interface based on integrated diffractive mirrors", 2016.

[Gis+16]   T. Gissibl, S. Thiele, A. Herkommer, and H. Giessen, "Sub-micrometre accurate free-form optics by three-dimensional printing on single-mode fibres", Nature Communications **7**, 11763 (2016).

[Giu+13]   M. Giustina, A. Mech, S. Ramelow, B. Wittmann, J. Kofler, J. Beyer, A. Lita, B. Calkins, T. Gerrits, S. W. Nam, R. Ursin, and A. Zeilinger, "Bell violation using entangled photons without the fair-sampling assumption", Nature **497**, 227–230 (2013).

[Giu+15]   M. Giustina, M. A. M. Versteegh, S. Wengerowsky, J. Handsteiner, A. Hochrainer, K. Phelan, F. Steinlechner, J. Kofler, J.-Å. Larsson, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, J. Beyer, T. Gerrits, A. E. Lita, L. K. Shalm, S. W. Nam, T. Scheidl, R. Ursin, B. Wittmann, and A. Zeilinger, "Significant-Loophole-Free Test of Bell's Theorem with Entangled Photons", Physical Review Letters **115**, 250401 (2015).

[GKG97]   M. Gu, P. C. Ke, and X. S. Gan, "Trapping force by a high numerical-aperture microscope objective obeying the sine condition", Review of Scientific Instruments **68**, 3666–3668 (1997).

[GLN05]   A. Gilchrist, N. K. Langford, and M. A. Nielsen, "Distance measures to compare real and ideal quantum processes", Physical Review A **71**, 1–14 (2005).

[Glo+15]   T. F. Gloger, P. Kaufmann, D. Kaufmann, M. T. Baig, T. Collath, M. Johanning, and C. Wunderlich, "Ion-trajectory analysis for micromotion minimization and the measurement of small forces", Physical Review A **92**, 043421 (2015).

[Gou+18]   É. Gouzien, B. Fedrici, A. Zavatta, S. Tanzilli, and V. D'Auria, "Quantum description of timing jitter for single-photon ON-OFF detectors", Physical Review A **98**, 013833 (2018).

[GSM20]   J. C. Garcia-Escartin, S. Sajeed, and V. Makarov, "Attacking quantum key distribution by light injection via ventilation openings", PLOS ONE **15**, edited by K.-K. R. Choo, e0236630 (2020).

[GVB16]   I. Gris-Sánchez, D. Van Ras, and T. A. Birks, "The Airy fiber: an optical fiber that guides light diffracted by a circular aperture", Optica **3**, 270 (2016).

[Ham00]   M. W. Hamilton, "Phase shifts in multilayer dielectric beam splitters", American Journal of Physics **68**, 186–191 (2000).

[Han+21] J. Hannegan, U. Saha, J. D. Siverns, J. Cassell, E. Waks, and Q. Quraishi, "C-band single photons from a trapped ion via two-stage frequency conversion", Applied Physics Letters **119**, 084001 (2021).

[Här+13] A. Härter, A. Krükow, A. Brunner, and J. Hecker Denschlag, "Minimization of ion micromotion using ultracold atomic probes", Applied Physics Letters **102**, 221115 (2013).

[Har+14] 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**, 2–6 (2014).

[Har+16] 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).

[Hen+15] B. Hensen, H. Bernien, A. E. Dréau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenberg, R. F. L. Vermeulen, R. N. Schouten, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminiau, and R. Hanson, "Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres", Nature **526**, 682–686 (2015).

[HH20]   D. Hurley-Smith and J. Hernandez-Castro, "Quantum Leap and Crash: Searching and Finding Bias in Quantum Random Number Generators", ACM Transactions on Privacy and Security **23**, 1–25 (2020).

[HH99]   M. Horodecki and P. Horodecki, "Reduction criterion of separability and limits for a class of distillation protocols", Physical Review A **59**, 4206–4216 (1999).

[Hig+21] G. Higgins, S. Salim, C. Zhang, H. Parke, F. Pokorny, and M. Hennrich, "Micromotion minimization using Ramsey interferometry", New Journal of Physics **23**, 123028 (2021).

[Hir+20]   H. Hirzler, T. Feldker, H. Fürst, N. V. Ewald, E. Trimby, R. S. Lous, J. D. Arias Espinoza, M. Mazzanti, J. Joger, and R. Gerritsma, "Experimental setup for studying an ultracold mixture of trapped Yb^+-^6Li", Physical Review A **102**, 033109 (2020).

[Ho+20]   M. Ho, P. Sekatski, E. Y.-Z. Tan, R. Renner, J.-D. Bancal, and N. Sangouard, "Noisy Preprocessing Facilitates a Photonic Realization of Device-Independent Quantum Key Distribution", Physical Review Letters **124**, 230502 (2020).

[Hof+12]   J. Hofmann, M. Krug, N. Ortegel, L. Gérard, M. Weber, W. Rosenfeld, and H. Weinfurter, "Heralded Entanglement Between Widely Separated Atoms", Science **337**, 72–75 (2012).

[Hom13]   J. P. Home, "Quantum Science and Metrology with Mixed-Species Ion Chains", Advances in Atomic, Molecular and Optical Physics **62**, 1–47 (2013).

[HOM87]   C. K. Hong, Z. Y. Ou, and L. Mandel, "Measurement of subpicosecond time intervals between two photons by interference", Physical Review Letters **59**, 2044–2046 (1987).

[Hor+09]   R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, "Quantum entanglement", Reviews of Modern Physics **81**, 865–942 (2009).

[Hou08]   M. G. House, "Analytic model for electrostatic fields in surface-electrode ion traps", Physical Review A **78**, 033402 (2008).

[Hub+14]   T. Huber, A. Lambrecht, J. Schmidt, L. Karpa, and T. Schaetz, "A far-off-resonance optical trap for a Ba^+ ion", Nature Communications **5**, 5587 (2014).

[Huc+15]   D. Hucul, I. V. Inlek, G. Vittorini, C. Crocker, S. Debnath, S. M. Clark, and C. Monroe, "Modular Entanglement of Atomic Qubits using both Photons and Phonons", Nature Physics **11**, 37–42 (2015).

[Hug+20]   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).

[Hug21]   A. C. Hughes, "Benchmarking memory and logic gates for trapped-ion quantum computing", DPhil thesis (University of Oxford, Nov. 16, 2021).

[Hum+18]   P. C. Humphreys, N. Kalb, J. P. J. Morits, R. N. Schouten, R. F. L. Vermeulen, D. J. Twitchen, M. Markham, and R. Hanson, "Deterministic delivery of remote entanglement on a quantum network", Nature **558**, 268–273 (2018).

[Hun+12]   N. Huntemann, M. Okhapkin, B. Lipphardt, S. Weyers, C. Tamm, and E. Peik, "High-accuracy optical clock based on the octupole transition in ^{171}Yb$^+$", Physical Review Letters **108**, 1–5 (2012).

[ITU11]   Y. Ibaraki, U. Tanaka, and S. Urabe, "Detection of parametric resonance of trapped ions for micromotion compensation", Applied Physics B **105**, 219–223 (2011).

[Iva22]   K. Ivanov, "Symmetry in design and decoding of polar-like codes", Doctoral thesis (École Polytechnique Fédérale de Lausanne, 2022).

[JC93]   J. Jin and D. A. Church, "Precision lifetimes for the Ca$^+$ $4p$ ^{2}P levels: Experiment challenges theory at the 1% level", Physical Review Letters **70**, 3213–3216 (1993).

[KA20]   M. Kessler and R. Arnon-Friedman, "Device-Independent Randomness Amplification and Privatization", IEEE Journal on Selected Areas in Information Theory **1**, 568–584 (2020).

[Kal+17]   N. Kalb, A. A. Reiserer, J. J. W. Bakermans, S. J. Kamerling, N. H. Nickerson, S. C. Benjamin, D. J. Twitchen, M. Markham, and R. Hanson, "Entanglement distillation between solid-state quantum network nodes", Science **356**, 928–932 (2017).

[Kas+23]   E. Kassa, W. Hughes, S. Gao, and J. F. Goodwin, "Effects of cavity birefringence in polarisation-encoded quantum networks", New Journal of Physics **25**, 013004 (2023).

[Kel+15]   J. Keller, H. L. Partner, T. Burgermeister, and T. E. Mehlstäubler, "Precise determination of micromotion for trapped-ion optical clocks", Journal of Applied Physics **118**, 104501 (2015).

[Kel22]   M. Keller, "Cavity-QED with single trapped ions", Contemporary Physics, 1–14 (2022).

[Kim08]   H. J. Kimble, "The Quantum Internet", Nature **453**, 1023–1030 (2008).

[Kin+12]   S. A. King, R. M. Godun, S. A. Webster, H. S. Margolis, L. A. M. Johnson, K. Szymaniec, P. E. G. Baird, and P. Gill, "Absolute frequency measurement of the ^{2}S$_{1/2}$ – ^{2}F$_{7/2}$ electric octupole transition in a single ion of ^{171}Yb$^+$ with 10^{-15} fractional uncertainty", New Journal of Physics **14**, 013045 (2012).

[KL15]   J. Katz and Y. Lindell, *Introduction to Modern Cryptography*, Second Edition, Chapman & Hall/CRC cryptography and network security (CRC Press/Taylor & Francis, Boca Raton, 2015), 583 pp.

[KM76]   H. J. Kimble and L. Mandel, "Theory of resonance fluorescence", Physical Review A **13**, 2123–2144 (1976).

[KMK11]   T. Kim, P. Maunz, and J. Kim, "Efficient collection of single photons emitted from a trapped ion into a single-mode fiber for scalable quantum-information processing", Physical Review A **84**, 1–10 (2011).

[KMW02]   D. Kielpinski, C. Monroe, and D. J. Wineland, "Architecture for a large-scale ion-trap quantum computer.", Nature **417**, 709–711 (2002).

[Kni+08]   E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, "Randomized benchmarking of quantum gates", Physical Review A **77**, 012307 (2008).

[Kot+14]   S. Kotler, N. Akerman, N. Navon, Y. Glickman, and R. Ozeri, "Measurement of the magnetic interaction between two bound electrons of two separate ions", Nature **510**, 376–380 (2014).

[Kro07]   T. Krovetz, *VMAC source code*, available at `https://www.fastcrypto.org/vmac`, 2007.

[Kru+17]   V. Krutyanskiy, M. Meraner, J. Schupp, and B. P. Lanyon, "Polarisation-preserving photon frequency conversion from a trapped-ion-compatible wavelength to the telecom C-band", Applied Physics B **123**, 228 (2017).

[Kru+19]   V. Krutyanskiy, M. Meraner, J. Schupp, V. Krcmarsky, H. Hainzer, and B. P. Lanyon, "Light-matter entanglement over 50 km of optical fibre", npj Quantum Information **5**, 72 (2019).

[Kru+23]   V. Krutyanskiy, M. Galli, V. Krcmarsky, S. Baier, D. A. Fioretto, Y. Pu, A. Mazloom, P. Sekatski, M. Canteri, M. Teller, J. Schupp, J. Bate, M. Meraner, N. Sangouard, B. P. Lanyon, and T. E. Northup, "Entanglement of Trapped-Ion Qubits Separated by 230 Meters", Physical Review Letters **130**, 050803 (2023).

[Lam52]   W. E. Lamb, "Fine Structure of the Hydrogen Atom. III", Physical Review **85**, 259–276 (1952).

[Lan+12]   H. Landa, M. Drewsen, B. Reznik, and A. Retzker, "Classical and quantum modes of coupled Mathieu equations", Journal of Physics A: Mathematical and Theoretical **45**, 455305 (2012).

[LCT14]   H.-K. Lo, M. Curty, and K. Tamaki, "Secure quantum key distribution", Nature Photonics **8**, 595–604 (2014).

[Leg+03]   T. Legero, T. Wilk, A. Kuhn, and G. Rempe, "Time-resolved two-photon quantum interference", Applied Physics B **77**, 797–802 (2003).

[Leg+06]   T. Legero, T. Wilk, A. Kuhn, and G. Rempe, "Characterization of Single Photons Using Two-Photon Interference", in *Advances In Atomic, Molecular, and Optical Physics*, Vol. 53 (Elsevier, 2006), pp. 253–289.

[Lei+03]   D. Leibfried, R. Blatt, C. Monroe, and D. Wineland, "Quantum dynamics of single trapped ions", Reviews of Modern Physics **75**, 281–324 (2003).

[Lei+07]   D. Leibfried, E. Knill, C. Ospelkaus, and D. J. Wineland, "Transport quantum logic gates for trapped ions", Physical Review A **76**, 32324 (2007).

[Lem+93]   V. E. Lembessis, M. Babiker, C. Baxter, and R. Loudon, "Theory of radiation forces and momenta for mobile atoms in light fields", Physical Review A **48**, 1594–1603 (1993).

[Leo+16]   V. Leong, M. A. Seidler, M. Steiner, A. Cerè, and C. Kurtsiefer, "Time-resolved scattering of a single photon by a single atom", Nature Communications **7**, 13716 (2016).

[Let+05]   V. Letchumanan, M. A. Wilson, P. Gill, and A. G. Sinclair, "Lifetime measurement of the metastable $4d^2 D_{5/2}$ Level in $^88Sr^+$ using a single trapped ion", Physical Review A **72**, 012509 (2005).

[Let+11]   M. Lettner, M. Mücke, S. Riedl, C. Vo, C. Hahn, S. Baur, J. Bochmann, S. Ritter, S. Dürr, and G. Rempe, "Remote Entanglement between a Single Atom and a Bose-Einstein Condensate", Physical Review Letters **106**, 210503 (2011).

[Lia+17]   S.-K. Liao, W.-Q. Cai, W.-Y. Liu, L. Zhang, Y. Li, J.-G. Ren, J. Yin, Q. Shen, Y. Cao, Z.-P. Li, F.-Z. Li, X.-W. Chen, L.-H. Sun, J.-J. Jia, J.-C. Wu, X.-J. Jiang, J.-F. Wang, Y.-M. Huang, Q. Wang, Y.-L. Zhou, L. Deng, T. Xi, L. Ma, T. Hu, Q. Zhang, Y.-A. Chen, N.-L. Liu, X.-B. Wang, Z.-C. Zhu, C.-Y. Lu, R. Shu, C.-Z. Peng, J.-Y. Wang, and J.-W. Pan, "Satellite-to-ground quantum key distribution", Nature **549**, 43–47 (2017).

[Liu+22]   W.-Z. Liu, Y.-Z. Zhang, Y.-Z. Zhen, M.-H. Li, Y. Liu, J. Fan, F. Xu, Q. Zhang, and J.-W. Pan, "Toward a Photonic Demonstration of Device-Independent Quantum Key Distribution", Physical Review Letters **129**, 050502 (2022).

[LMC05]   H.-K. Lo, X. Ma, and K. Chen, "Decoy State Quantum Key Distribution", Physical Review Letters **94**, 230504 (2005).

[Loh+13]   H. Loh, K. C. Cossel, M. C. Grau, K. K. Ni, E. R. Meyer, J. L. Bohn, J. Ye, and E. A. Cornell, "Precision Spectroscopy of Polarized Molecules in an Ion Trap", Science **342**, 1220–1222 (2013).

[Lou00]   R. Loudon, *The Quantum Theory of Light*, 3rd ed., Oxford Science Publications (Oxford University Press, 2000), 454 pp.

[LS07]   P. L'Ecuyer and R. Simard, "TestU01: A C Library for Empirical Testing of Random Number Generators", ACM Trans. Math. Softw. **33**, 10 . 1145 / 1268776 . 1268777 (2007).

[Luc+04]   D. Lucas, A. Ramos, J. Home, M. McDonnell, S. Nakayama, J.-P. Stacey, S. Webster, D. Stacey, and A. Steane, "Isotope-selective photoionization for calcium ion trapping", Physical Review A **69**, 012711 (2004).

[Luo+09]   L. Luo, D. Hayes, T. A. Manning, D. N. Matsukevich, P. Maunz, S. Olmschenk, J. D. Sterk, and C. Monroe, "Protocols and techniques for a scalable atom-photon quantum network", Fortschritte der Physik **57**, 1133–1152 (2009).

[Lyd+10]   L. Lydersen, C. Wiechers, C. Wittmann, D. Elser, J. Skaar, and V. Makarov, "Hacking commercial quantum cryptography systems by tailored bright illumination", Nature Photonics **4**, 686–689 (2010).

[Mai+01]   A. Mair, A. Vaziri, G. Weihs, and A. Zeilinger, "Entanglement of the orbital angular momentum states of photons", Nature **412**, 313–316 (2001).

[Man+99]   S. Mannervik, J. Lidberg, L.-O. Norlin, P. Royen, A. Schmitt, W. Shi, and X. Tordoir, "Lifetime Measurement of the Metastable $4d^2$ $D_{3/2}$ Level in Sr^+ by Optical Pumping of a Stored Ion Beam", Physical Review Letters **83**, 698–701 (1999).

[Mat+08]   D. N. Matsukevich, P. Maunz, D. L. Moehring, S. Olmschenk, and C. Monroe, "Bell Inequality Violation with Two Remote Atomic Qubits", Physical Review Letters **100**, 150404 (2008).

[Mat+21]   C. Matthiesen, Q. Yu, J. Guo, A. M. Alonso, and H. Häffner, "Trapping Electrons in a Room-Temperature Microwave Paul Trap", Physical Review X **11**, 011019 (2021).

[Mau+09]   P. Maunz, S. Olmschenk, D. Hayes, D. N. Matsukevich, L.-M. Duan, and C. Monroe, "Heralded Quantum Gate between Remote Quantum Memories", Physical Review Letters **102**, 250502 (2009).

[Mau11]   U. Maurer, "Constructive cryptography–a new paradigm for security definitions and proofs", in Joint Workshop on Theory of Security and Applications (Springer, 2011), pp. 33–56.

[Mau14]   W. Mauerer, *Trevisan's extractor source code*, available at `https://github.com/wolfgangmauerer/libtrevisan`, 2014.

[Mau16]   P. L. W. Maunz, *High Optical Access Trap 2.0*, Technical Report SAND-2016-0796R (Sandia National Laboratories, Albuquerque, New Mexico, Jan. 26, 2016).

[MB08]   J. Metz and S. D. Barrett, "Effect of frequency-mismatched photons in quantum-information processing", Physical Review A **77**, 042323 (2008).

[Meh+16]   K. K. Mehta, C. D. Bruzewicz, R. McConnell, R. J. Ram, J. M. Sage, and J. Chiaverini, "Integrated optical addressing of an ion qubit", Nature Nanotechnology **11**, 1066–1070 (2016).

[Meh+20] K. K. Mehta, C. Zhang, M. Malinowski, T.-L. Nguyen, M. Stadler, and J. P. Home, "Integrated optical multi-ion quantum logic", Nature **586**, 533–537 (2020).

[Mei+16] Z. Meir, T. Sikorsky, R. Ben-shlomi, N. Akerman, Y. Dallal, and R. Ozeri, "Dynamics of a Ground-State Cooled Ion Colliding with Ultracold Atoms", Physical Review Letters **117**, 243401 (2016).

[Mei+18] Z. Meir, T. Sikorsky, R. Ben-shlomi, N. Akerman, M. Pinkas, Y. Dallal, and R. Ozeri, "Experimental apparatus for overlapping a ground-state cooled ion with ultracold atoms", Journal of Modern Optics **65**, 501–519 (2018).

[Mer+19] B. Merkel, K. Thirumalai, J. E. Tarlton, V. M. Schäfer, C. J. Ballance, T. P. Harty, and D. M. Lucas, "Magnetic field stabilization system for atomic physics experiments", Review of Scientific Instruments **90**, 044702 (2019).

[Mer+20] M. Meraner, A. Mazloom, V. Krutyanskiy, V. Krcmarsky, J. Schupp, D. A. Fioretto, P. Sekatski, T. E. Northup, N. Sangouard, and B. P. Lanyon, "Indistinguishable photons from a trapped-ion quantum network node", Physical Review A **102**, 052614 (2020).

[Mer78] R. C. Merkle, "Secure communications over insecure channels", Communications of the ACM **21**, 294–299 (1978).

[MF11] M. Mayeh and F. Farahi, "Laser beam shaping and mode conversion in optical fibers", Photonic Sensors **1**, 187–198 (2011).

[MF13] T. Morimae and K. Fujii, "Blind quantum computation protocol in which Alice only makes measurements", Physical Review A **87**, 1–5 (2013).

[Mil76] P. Milonni, "Semiclassical and quantum-electrodynamical approaches in nonrelativistic radiation theory", Physics Reports **25**, 1–81 (1976).

[Mil94] P. W. Milonni, *The quantum vacuum: an introduction to quantum electrodynamics* (Academic Press, Boston, 1994), 522 pp.

[MK13] C. Monroe and J. Kim, "Scaling the ion trap quantum processor", Science **339**, 1164–9 (2013).

[Moe+07] D. L. Moehring, P. Maunz, S. Olmschenk, K. C. Younge, D. N. Matsukevich, L.-M. Duan, and C. Monroe, "Entanglement of single-atom quantum bits at a distance", Nature **449**, 68–71 (2007).

[Moh+19] A. Mohammadi, J. Wolf, A. Krükow, M. Deiß, and J. Hecker Denschlag, "Minimizing rf-induced excess micromotion of a trapped ion with the help of ultracold atoms", Applied Physics B **125**, 122 (2019).

[Mon+14]   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**, 1–16 (2014).

[Moo+22]   G. Moody, V. J. Sorger, D. J. Blumenthal, P. W. Juodawlkis, W. Loh, C. Sorace-Agaskar, A. E. Jones, K. C. Balram, J. C. F. Matthews, A. Laing, M. Davanco, L. Chang, J. E. Bowers, N. Quack, C. Galland, I. Aharonovich, M. A. Wolff, C. Schuck, N. Sinclair, M. Lončar, T. Komljenovic, D. Weld, S. Mookherjea, S. Buckley, M. Radulaski, S. Reitzenstein, B. Pingault, B. Machielse, D. Mukhopadhyay, A. Akimov, A. Zheltikov, G. S. Agarwal, K. Srinivasan, J. Lu, H. X. Tang, W. Jiang, T. P. McKenna, A. H. Safavi-Naeini, S. Steinhauer, A. W. Elshaari, V. Zwiller, P. S. Davids, N. Martinez, M. Gehl, J. Chiaverini, K. K. Mehta, J. Romero, N. B. Lingaraju, A. M. Weiner, D. Peace, R. Cernansky, M. Lobino, E. Diamanti, L. T. Vidarte, and R. M. Camacho, "2022 Roadmap on integrated quantum photonics", Journal of Physics: Photonics **4**, 012501 (2022).

[MSZ05]   A. Muthukrishnan, M. O. Scully, and M. S. Zubairy, "The photon wave function", in , edited by C. Roychoudhuri and K. Creath (Aug. 1, 2005), p. 287.

[Mül+17]   P. Müller, T. Tentrup, M. Bienert, G. Morigi, and J. Eschner, "Spectral properties of single photons from quantum emitters", Physical Review A **96**, 023861 (2017).

[Mur+19]   G. Murta, S. B. van Dam, J. Ribeiro, R. Hanson, and S. Wehner, "Towards a realization of device-independent quantum key distribution", Quantum Science and Technology **4**, 035011 (2019).

[MY04]   D. Mayers and A. Yao, "Self testing quantum apparatus", Quantum Information & Computation **4**, 273–286 (2004).

[Mye+08]   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).

[Nad+21]   D. P. Nadlinger, P. Drmota, D. Main, B. C. Nichol, G. Araneda, R. Srinivas, L. J. Stephenson, C. J. Ballance, and D. M. Lucas, "Micromotion minimisation by synchronous detection of parametrically excited motion", preprint **arXiv:2107.00056** (2021).

[Nad+22]   D. P. Nadlinger, P. Drmota, B. C. Nichol, G. Araneda, D. Main, R. Srinivas, D. M. Lucas, C. J. Ballance, K. Ivanov, E. Y.-Z. Tan, P. Sekatski, R. L. Urbanke,

R. Renner, N. Sangouard, and J.-D. Bancal, "Experimental quantum key distribution certified by Bell's theorem", Nature **607**, 682–686 (2022).

[Nad16]   D. Nadlinger, "Entanglement between Trapped Strontium Ions and Photons", MSc Thesis (ETH Zürich, 2016).

[Nag+02]   K. Nagayama, M. Kakui, M. Matsui, T. Saitoh, and Y. Chigusa, "Ultra-low-loss (0.1484 dB / km) pure silica core fibre and extension of transmission distance", Electronics Letters **38**, 1168 (2002).

[Nar+11]   S. Narayanan, N. Daniilidis, S. A. Möller, R. Clark, F. Ziesel, K. Singer, F. Schmidt-Kaler, and H. Häffner, "Electric field compensation and sensing with a single ion in a planar trap", Journal of Applied Physics **110**, 114909 (2011).

[NC11]   M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (2011).

[NFB14]   N. H. Nickerson, J. F. Fitzsimons, and S. C. Benjamin, "Freely scalable quantum technologies using cells of 5-to-50 qubits with very lossy and noisy photonic links", Physical Review X **4**, 1–17 (2014).

[NH06]   L. Novotny and B. Hecht, *Principles of Nano-Optics* (Cambridge University Press, 2006).

[Nic+22]   B. C. Nichol, R. Srinivas, D. P. Nadlinger, P. Drmota, D. Main, G. Araneda, C. J. Ballance, and D. M. Lucas, "An elementary quantum network of entangled optical atomic clocks", Nature **609**, 689–694 (2022).

[Nic22]   B. C. Nichol, "An elementary quantum network of entangled optical atomic clocks", DPhil thesis (University of Oxford, 2022).

[Nie+21]   D. Niemietz, P. Farrera, S. Langenfeld, and G. Rempe, "Nondestructive detection of photonic qubits", Nature **591**, 570–574 (2021).

[Nig+16]   R. Nigmatullin, C. J. Ballance, N. De Beaudrap, and S. C. Benjamin, "Minimally complex ion traps as modules for quantum communication and computing", New Journal of Physics **18**, 103028 (2016).

[Nöl+13]   C. Nölleke, A. Neuzner, A. Reiserer, C. Hahn, G. Rempe, and S. Ritter, "Efficient Teleportation Between Remote Single-Atom Quantum Memories", Physical Review Letters **110**, 140403 (2013).

[NTH12]   C. M. Natarajan, M. G. Tanner, and R. H. Hadfield, "Superconducting nanowire single-photon detectors: physics and applications", Superconductor Science and Technology **25**, 063001 (2012).

[Pau90]   W. Paul, "Electromagnetic traps for charged and neutral particles", Reviews of Modern Physics **62**, 531–540 (1990).

[PBL95]   E. H. Pinnington, R. W. Berends, and M. Lumsden, "Studies of laser-induced fluorescence in fast beams of Sr$^+$ and Ba$^+$ ions", Journal of Physics B: At. Mol. Opt. Phys. **28**, 2095 (1995).

[Pel+11]   J. S. Pelc, L. Ma, C. R. Phillips, Q. Zhang, C. Langrock, O. Slattery, X. Tang, and M. M. Fejer, "Long-wavelength-pumped upconversion single-photon detector at 1550 nm: performance and noise analysis", Optics Express **19**, 21445–21456 (2011).

[Pet+22]   M. Petrov, I. Radchenko, D. Steiger, R. Renner, M. Troyer, and V. Makarov, "Independent quality assessment of a commercial quantum random number generator", EPJ Quantum Technology **9**, 17 (2022).

[Pfa+14]   W. Pfaff, B. J. Hensen, H. Bernien, S. B. van Dam, M. S. Blok, T. H. Taminiau, M. J. Tiggelman, R. N. Schouten, M. Markham, D. J. Twitchen, and R. Hanson, "Unconditional quantum teleportation between distant solid-state quantum bits", Science **345**, 532–535 (2014).

[Pir+09]   S. Pironio, A. Acín, N. Brunner, N. Gisin, S. Massar, and V. Scarani, "Device-independent quantum key distribution secure against collective attacks", New Journal of Physics **11**, 045021 (2009).

[Pir+10]   S. Pironio, A. Acín, S. Massar, A. B. de la Giroday, D. N. Matsukevich, P. Maunz, S. Olmschenk, D. Hayes, L. Luo, T. A. Manning, and C. Monroe, "Random numbers certified by Bell's theorem", Nature **464**, 1021–1024 (2010).

[Pir+20]   S. Pirandola, U. L. Andersen, L. Banchi, M. Berta, D. Bunandar, R. Colbeck, D. Englund, T. Gehring, C. Lupo, C. Ottaviani, J. L. Pereira, M. Razavi, J. Shamsul Shaari, M. Tomamichel, V. C. Usenko, G. Vallone, P. Villoresi, and P. Wallden, "Advances in quantum cryptography", Advances in Optics and Photonics **12**, 1012 (2020).

[PK98]   M. B. Plenio and P. L. Knight, "The quantum-jump approach to dissipative dynamics in quantum optics", Reviews of Modern Physics **70**, 101–144 (1998).

[Ple05]   M. B. Plenio, "Logarithmic Negativity: A Full Entanglement Monotone That is not Convex", Physical Review Letters **95**, 090503 (2005).

[Pom+21]   M. Pompili, S. L. N. Hermans, S. Baier, H. K. C. Beukers, P. C. Humphreys, R. N. Schouten, R. F. L. Vermeulen, M. J. Tiggelman, L. dos Santos Martins, B. Dirkse, S. Wehner, and R. Hanson, "Realization of a multinode quantum network of remote solid-state qubits", Science **372**, 259–264 (2021).

[PR22]   C. Portmann and R. Renner, "Security in quantum cryptography", Rev. Mod. Phys. **94**, 56 (2022).

[PR94]   S. Popescu and D. Rohrlich, "Quantum nonlocality as an axiom", Foundations of Physics **24**, 379–385 (1994).

[Pre18]   J. Preskill, "Quantum Computing in the NISQ era and beyond", Quantum **2**, 79 (2018).

[PZ59]   E. A. Power and S. Zienau, "Coulomb gauge in non-relativistic quantum electro-dynamics and the shape of spectral lines", Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences **251**, 427–454 (1959).

[Reh+07]   J. Reháček, Z. Hradil, E. Knill, and A. I. Lvovsky, "Diluted maximum-likelihood algorithm for quantum tomography", Physical Review A **75**, 042108 (2007).

[RF17]   E. Rousseau and D. Felbacq, "The quantum-optics Hamiltonian in the Multipolar gauge", Scientific Reports **7**, 11115 (2017).

[Rie+18]   R. Riedinger, A. Wallucks, I. Marinković, C. Löschnauer, M. Aspelmeyer, S. Hong, and S. Gröblacher, "Remote quantum entanglement between two micromechanical oscillators", Nature **556**, 473–477 (2018).

[Rit+12]   S. Ritter, C. Nölleke, C. Hahn, A. Reiserer, A. Neuzner, M. Uphoff, M. Mücke, E. Figueroa, J. Bochmann, and G. Rempe, "An elementary quantum network of single atoms in optical cavities", Nature **484**, 195–200 (2012).

[Ros+17]   W. Rosenfeld, D. Burchardt, R. Garthoff, K. Redeker, N. Ortegel, M. Rau, and H. Weinfurter, "Event-Ready Bell Test Using Entangled Atoms Simultaneously Closing Detection and Locality Loopholes", Physical Review Letters **119**, 010402 (2017).

[RR15]   A. Reiserer and G. Rempe, "Cavity-based quantum networks with single atoms and optical photons", Reviews of Modern Physics **87**, 1379–1418 (2015).

[RS16]   S. Rajasekar and M. A. F. Sanjuan, "Parametric Resonance", in *Nonlinear Resonances*, Springer Series in Synergetics (Springer International Publishing, Cham, 2016).

[RSA78]   R. L. Rivest, A. Shamir, and L. Adleman, "A method for obtaining digital signatures and public-key cryptosystems", Communications of the ACM **21**, 120–126 (1978).

[Ruf+21]   M. Ruf, N. H. Wan, H. Choi, D. Englund, and R. Hanson, "Quantum networks based on color centers in diamond", Journal of Applied Physics **130**, 070901 (2021).

[RW59] B. Richards and E. Wolf, "Electromagnetic diffraction in optical systems, II. Structure of the image field in an aplanatic system", Proc. R. Soc. Lond. A **253**, 358–379 (1959).

[RZ92] K. Rzazewski and W. Zakowicz, "Spontaneous emission from an extended wavepacket", Journal of Physics B: Atomic, Molecular and Optical Physics **25**, L319–L322 (1992).

[Sah+06] B. K. Sahoo, Md. R. Islam, B. P. Das, R. K. Chaudhuri, and D. Mukherjee, "Lifetimes of the metastable $^2D_{3/2,5/2}$ states in Ca$^+$, Sr$^+$, and Ba$^+$", Physical Review A **74**, 062504 (2006).

[Sak+20] H. Sakr, Y. Chen, G. T. Jasion, T. D. Bradley, J. R. Hayes, H. C. H. Mulvad, I. A. Davidson, E. Numkam Fokoua, and F. Poletti, "Hollow core optical fibres with comparable attenuation to silica fibres between 600 and 1100 nm", Nature Communications **11**, 6030 (2020).

[San12] J. E. Sansonetti, "Wavelengths, Transition Probabilities, and Energy Levels for the Spectra of Strontium Ions (Sr II through Sr XXXVIII)", Journal of Physical and Chemical Reference Data **41** (2012).

[Sca+09] V. Scarani, H. Bechmann-Pasquinucci, N. J. Cerf, M. Dušek, N. Lütkenhaus, and M. Peev, "The security of practical quantum key distribution", Reviews of Modern Physics **81**, 1301–1350 (2009).

[Sch+05] P. O. Schmidt, T. Rosenband, C. Langer, W. M. Itano, J. C. Bergquist, and D. J. Wineland, "Spectroscopy Using Quantum Logic", Science **309**, 749–752 (2005).

[Sch+07] T. Schmitt-Manderbach, H. Weier, M. Fürst, R. Ursin, F. Tiefenbacher, T. Scheidl, J. Perdigues, Z. Sodnik, C. Kurtsiefer, J. G. Rarity, A. Zeilinger, and H. Weinfurter, "Experimental Demonstration of Free-Space Decoy-State Quantum Key Distribution over 144 km", Physical Review Letters **98**, 010504 (2007).

[Sch+20] J. Schmidt, D. Hönig, P. Weckesser, F. Thielemann, T. Schaetz, and L. Karpa, "Mass-selective removal of ions from Paul traps using parametric excitation", Applied Physics B **126**, 176 (2020).

[Sch+21a] J. Schupp, V. Krcmarsky, V. Krutyanskiy, M. Meraner, T. Northup, and B. Lanyon, "Interface between Trapped-Ion Qubits and Traveling Photons with Close-to-Optimal Efficiency", PRX Quantum **2**, 020331 (2021).

[Sch+21b] R. Schwonnek, K. T. Goh, I. W. Primaatmaja, E. Y.-Z. Tan, R. Wolf, V. Scarani, and C. C.-W. Lim, "Device-independent quantum key distribution with random key basis", Nature Communications **12**, 2880 (2021).

[SD06]   C. Schwartz and A. Dogariu, "Conservation of angular momentum of light in single scattering", Optics Express **14**, 8425 (2006).

[Sek+21]   P. Sekatski, J.-D. Bancal, X. Valcarce, E. Y.-Z. Tan, R. Renner, and N. Sangouard, "Device-independent quantum key distribution from generalized CHSH inequalities", Quantum **5**, 444 (2021).

[Sep+19]   M. A. Sepiol, A. C. Hughes, J. E. Tarlton, D. P. Nadlinger, T. G. Ballance, C. J. Ballance, T. P. Harty, A. M. Steane, J. F. Goodwin, and D. M. Lucas, "Probing Qubit Memory Errors at the Part-per-Million Level", Physical Review Letters **123**, 110503 (2019).

[SG93]   C. Sheppard and M. Gu, "Imaging by a High Aperture Optical System", Journal of Modern Optics **40**, 1631–1651 (1993).

[SGV17]   G. B. Silva, S. Glancy, and H. M. Vasconcelos, "Investigating bias in maximum-likelihood quantum-state tomography", Physical Review A **95**, 022107 (2017).

[Sha+15]   L. K. Shalm, E. Meyer-Scott, B. G. Christensen, P. Bierhorst, M. A. Wayne, M. J. Stevens, T. Gerrits, S. Glancy, D. R. Hamel, M. S. Allman, K. J. Coakley, S. D. Dyer, C. Hodge, A. E. Lita, V. B. Verma, C. Lambrocco, E. Tortorici, A. L. Migdall, Y. Zhang, D. R. Kumor, W. H. Farr, F. Marsili, M. D. Shaw, J. A. Stern, C. Abellán, W. Amaya, V. Pruneri, T. Jennewein, M. W. Mitchell, P. G. Kwiat, J. C. Bienfang, R. P. Mirin, E. Knill, and S. W. Nam, "Strong Loophole-Free Test of Local Realism", Physical Review Letters **115**, 250402 (2015).

[Sho94]   P. Shor, "Algorithms for quantum computation: discrete logarithms and factoring", in Proceedings 35th Annual Symposium on Foundations of Computer Science (1994), pp. 124–134.

[Sho95]   P. W. Shor, "Scheme for reducing decoherence in quantum computer memory", Physical Review A **52**, R2493–R2496 (1995).

[Shu+10]   G. Shu, N. Kurz, M. R. Dietrich, and B. B. Blinov, "Efficient fluorescence collection from trapped ions with an integrated spherical mirror", Physical Review A **81**, 042321 (2010).

[Sli+16]   D. H. Slichter, V. B. Verma, D. Leibfried, R. P. Mirin, S. W. Nam, and D. J. Wineland, "UV-sensitive superconducting nanowire single photon detectors for integration in an ion trap", 2016.

[Slo+13]   L. Slodička, G. Hétet, N. Röck, P. Schindler, M. Hennrich, and R. Blatt, "Atom-Atom Entanglement by Single-Photon Detection", Physical Review Letters **110**, 083603 (2013).

[SR07]   B. J. Smith and M. G. Raymer, "Photon wave functions, wave-packet quantization of light, and coherence theory", New Journal of Physics **9**, 414–414 (2007).

[Sri+14]   B. Srivathsan, G. K. Gulati, A. Cerè, B. Chng, and C. Kurtsiefer, "Reversing the Temporal Envelope of a Heralded Single Photon using a Cavity", Physical Review Letters **113**, 163601 (2014).

[Sri+21]   R. Srinivas, S. C. Burd, H. M. Knaack, R. T. Sutherland, A. Kwiatkowski, S. Glancy, E. Knill, D. J. Wineland, D. Leibfried, A. C. Wilson, D. T. C. Allcock, and D. H. Slichter, "High-fidelity laser-free universal control of trapped ion qubits", Nature **597**, 209–213 (2021).

[ST11]   J. Slane and S. Tragesser, "Analysis of Periodic Nonautonomous Inhomogeneous Systems", Nonlinear Dynamics and Systems Theory **11** (2011).

[Ste+20]   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).

[Ste19]   L. J. Stephenson, "Entanglement between nodes of a quantum network", DPhil thesis (University of Oxford, 2019).

[Ste96]   A. M. Steane, "Error Correcting Codes in Quantum Theory", Physical Review Letters **77**, 793–797 (1996).

[Sto+17]   R. Stockill, M. J. Stanley, L. Huthmacher, E. Clarke, M. Hugues, A. J. Miller, C. Matthiesen, C. Le Gall, and M. Atatüre, "Phase-Tuned Entangled State Generation between Distant Spin Qubits", Physical Review Letters **119**, 010503 (2017).

[Sto11]   E. Stock, "Self-organized Quantum Dots for Single Photon Sources", Doctoral thesis (Technische Universität Berlin, Feb. 8, 2011).

[SW82]   C. Sheppard and T. Wilson, "The image of a single point in microscopes of large numerical aperture", Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences **379**, 145–158 (1982).

[SZ97]   M. O. Scully and M. S. Zubairy, *Quantum Optics* (Cambridge University Press, 1997).

[Tam+18]   Y. Tamura, H. Sakuma, K. Morita, M. Suzuki, Y. Yamamoto, K. Shimada, Y. Honma, K. Sohma, T. Fujii, and T. Hasegawa, "The First 0.14-dB/km Loss Optical Fiber and its Impact on Submarine Transmission", Journal of Lightwave Technology **36**, 44–49 (2018).

[Tan+12]   U. Tanaka, K. Masuda, Y. Akimoto, K. Koda, Y. Ibaraki, and S. Urabe, "Micromotion compensation in a surface electrode trap by parametric excitation of trapped ions", Applied Physics B **107**, 907–912 (2012).

[Tan+20]   E. Y.-Z. Tan, P. Sekatski, J.-D. Bancal, R. Schwonnek, R. Renner, N. Sangouard, and C. C.-W. Lim, "Improved DIQKD protocols with finite-size analysis", Dec. 15, 2020.

[Tel+21]   M. Teller, D. A. Fioretto, P. C. Holz, P. Schindler, V. Messerer, K. Schüppert, Y. Zou, R. Blatt, J. Chiaverini, J. Sage, and T. E. Northup, "Heating of a Trapped Ion Induced by Dielectric Materials", Physical Review Letters **126**, 230505 (2021).

[TG66]   U. M. Titulaer and R. J. Glauber, "Density Operators for Coherent Fields", Physical Review **145**, 1041–1050 (1966).

[Thi19]   K. Thirumalai, "High-Fidelity Mixed Species Entanglement of Trapped Ions", DPhil thesis (University of Oxford, 2019).

[TLR20]   E. Y.-Z. Tan, C. C.-W. Lim, and R. Renner, "Advantage Distillation for Device-Independent Quantum Key Distribution", Physical Review Letters **124**, 020502 (2020).

[TMM21]   O. F. Thomas, W. McCutcheon, and D. P. S. McCutcheon, "A general framework for multimode Gaussian quantum optics and photo-detection: Application to Hong–Ou–Mandel interference with filtered heralded single photon sources", APL Photonics **6**, 040801 (2021).

[Ton09]   B. Toner, "Monogamy of non-local quantum correlations", Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences **465**, 59–69 (2009).

[Tre+09]   A. Treiber, A. Poppe, M. Hentschel, D. Ferrini, T. Lorünser, E. Querasser, T. Matyus, H. Hübel, and A. Zeilinger, "A fully automated entanglement-based quantum cryptography system for telecom fiber networks", New Journal of Physics **11**, 045013 (2009).

[Tre01]   L. Trevisan, "Extractors and Pseudorandom Generators", Journal of the ACM **48**, 860–879 (2001).

[Val+14]   R. Valivarthi, I. Lucio-Martinez, A. Rubenok, P. Chan, F. Marsili, V. B. Verma, M. D. Shaw, J. A. Stern, J. A. Slater, D. Oblak, S. W. Nam, and W. Tittel, "Efficient Bell state analyzer for time-bin qubits with fast-recovery WSi superconducting single photon detectors.", Optics express **22**, 24497–506 (2014).

[Vit+14]   G. Vittorini, D. Hucul, I. V. Inlek, C. Crocker, and C. Monroe, "Entanglement of distinguishable quantum memories", Physical Review A **90**, 040302 (2014).

[vLee+22] T. van Leent, M. Bock, F. Fertig, R. Garthoff, S. Eppelt, Y. Zhou, P. Malik, M. Seubert, T. Bauer, W. Rosenfeld, W. Zhang, C. Becher, and H. Weinfurter, "Entangling single atoms over 33 km telecom fibre", Nature **607**, 69–73 (2022).

[VV14] U. Vazirani and T. Vidick, "Fully device independent quantum key distribution", Physical Review Letters **113**, 140501 (2014).

[VW02] G. Vidal and R. F. Werner, "Computable measure of entanglement", Physical Review A **65**, 1–11 (2002).

[Wal+20] T. Walker, S. V. Kashanian, T. Ward, and M. Keller, "Improving the indistinguishability of single photons from an ion-cavity system", Physical Review A **102**, 032616 (2020).

[Wan+11] Y. Wang, J. Minář, L. Sheridan, and V. Scarani, "Efficient excitation of a two-level atom by a single photon in a propagating mode", Physical Review A **83**, 063842 (2011).

[Wan+21] 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).

[War+13] U. Warring, C. Ospelkaus, Y. Colombe, K. R. Brown, J. M. Amini, M. Carsjens, D. Leibfried, and D. J. Wineland, "Techniques for microwave near-field quantum control of trapped ions", Physical Review A **87**, 013437 (2013).

[WEH18] S. Wehner, D. Elkouss, and R. Hanson, "Quantum internet: A vision for the road ahead", Science **362**, eaam9288 (2018).

[Wei+11] H. Weier, H. Krauss, M. Rau, M. Fürst, S. Nauerth, and H. Weinfurter, "Quantum eavesdropping without interception: an attack exploiting the dead time of single-photon detectors", New Journal of Physics **13**, 073024 (2011).

[WGW95] W.-T. Wu, J. Griffin, and J. Wickert, "Perturbation Method for the Floquet Eigenvalues and Stability Boundary of Periodic Linear Systems", Journal of Sound and Vibration **182**, 245–257 (1995).

[Wil94] M. Wilkens, "Significance of Röntgen current in quantum optics: Spontaneous emission of moving atoms", Physical Review A **49**, 570–573 (1994).

[Win+98] D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. Meekhof, "Experimental issues in coherent quantum-state manipulation of trapped atomic ions", Journal of Research of the National Institute of Standards and Technology **103**, 259–328 (1998).

[Wol59] E. Wolf, "Electromagnetic diffraction in optical systems - I. An integral representation of the image field", Proc. R. Soc. Lond. A **253**, 349–357 (1959).

[Won+16]   J. D. Wong-Campos, K. G. Johnson, B. Neyenhuis, J. Mizrahi, and C. Monroe, "High-resolution adaptive imaging of a single atom", Nature Photonics **10**, 606–610 (2016).

[Woo20]   R. G. Woolley, "Power-Zienau-Woolley representations of nonrelativistic QED for atoms and molecules", Physical Review Research **2**, 013206 (2020).

[Woo71]   R. G. Woolley, "Molecular quantum electrodynamics", Proc. R. Soc. Lond. A **321**, 557–572 (1971).

[Woo98]   W. K. Wootters, "Entanglement of Formation of an Arbitrary State of Two Qubits", Physical Review Letters **80**, 2245–2248 (1998).

[Wri+18]   T. A. Wright, R. J. A. Francis-Jones, C. B. E. Gawith, J. N. Becker, P. M. Ledingham, P. G. R. Smith, J. Nunn, P. J. Mosley, B. Brecht, and I. A. Walmsley, "Two-Way Photonic Interface for Linking the Sr + Transition at 422 nm to the Telecommunication C Band", Physical Review Applied **10**, 044012 (2018).

[WW27]   V. Weisskopf and E. Wigner, "Berechnung der natürlichen Linienbreite auf Grund der Dirac'schen Lichttheorie", Proc. R. Soc. Lond. A **114**, 710 (1927).

[WWM02]   P. Warszawski, H. M. Wiseman, and H. Mabuchi, "Quantum trajectories for realistic detection", Physical Review A **65**, 023802 (2002).

[YDN93]   N. Yu, H. Dehmelt, and W. Nagourney, "Noise-free parametric energy multiplication for frequency measurements of an anharmonic mono-ion oscillator", Journal of Applied Physics **73**, 8650–8652 (1993).

[Zha+08]   Y. Zhao, C.-H. F. Fung, B. Qi, C. Chen, and H.-K. Lo, "Quantum hacking: Experimental demonstration of time-shift attack against practical quantum-key-distribution systems", Physical Review A **78**, 042333 (2008).

[Zha+14]   T.-M. Zhao, H. Zhang, J. Yang, Z.-R. Sang, X. Jiang, X.-H. Bao, and J.-W. Pan, "Entangling Different-Color Photons via Time-Resolved Measurement and Active Feed Forward", Physical Review Letters **112**, 103602 (2014).

[Zha+20]   C. Zhang, F. Pokorny, W. Li, G. Higgins, A. Pöschl, I. Lesanovsky, and M. Hennrich, "Submicrosecond entangling gate between trapped ions via Rydberg interaction", Nature **580**, 345–349 (2020).

[Zha+22]   W. Zhang, T. van Leent, K. Redeker, R. Garthoff, R. Schwonnek, F. Fertig, S. Eppelt, V. Scarani, C. C.-W. Lim, and H. Weinfurter, "A device-independent quantum key distribution system for distant users", Nature **607**, 687–691 (2022).

[Żuk+93]   M. Żukowski, A. Zeilinger, M. A. Horne, and A. K. Ekert, ""Event-ready-detectors" Bell experiment via entanglement swapping", Physical Review Letters **71**, 4287–4290 (1993).