

# **Unitary and Dissipative Trapped-Ion Entanglement Using Integrated Optics**

A thesis submitted to attain the degree of  
DOCTOR OF SCIENCES of ETH ZURICH  
(Dr. sc. ETH Zurich)

presented by

MACIEJ MALINOWSKI

M. Phys., University of Oxford  
born on 23.10.1992  
citizen of Poland

accepted on the recommendation of

Prof. Dr. J. P. Home  
Prof. Dr. H. Häffner



# Abstract

Trapped atomic ions are currently among the most advanced platforms for quantum information processing. While the individual elements of a universal quantum computer have all been demonstrated, it is an outstanding challenge to increase the system size to where it becomes computationally useful.

In this thesis, we describe the work towards building scalable trapped-ion quantum information processors. We discuss the construction and operation of a cryogenic (4K) vacuum system designed for testing novel ion traps. We present the laser setups used for cooling, manipulating and measuring the ions. Particular emphasis is placed on developing a narrow-linewidth high-power qubit control laser.

The ion trap at the heart of the cryogenic system is a novel surface-electrode trap with integrated photonics. In contrast to previous demonstrations, where all the lasers were delivered free-space, the qubit control light is delivered through optical fibres, coupled into trap-integrated waveguides, and focused onto the ions using dielectric grating couplers. This allows us to reduce beam-pointing fluctuations, and create small spot sizes which result in high laser-ion coupling rates.

We extensively characterise multiple traps with integrated photonics, finding several challenges associated with material charging and heating rates. Nonetheless, we demonstrate high-quality quantum control using trap-integrated light. Notably, we implement two-qubit gates based on the bichromatic Mølmer-Sørensen interaction. We use them to prepare maximally-entangled states of two  $^{40}\text{Ca}^+$  ions with 99.45(13)% fidelity, limited by well-characterised error sources.

In addition to laser-based control of the optical qubit, we use global radio-frequency magnetic fields to manipulate the Zeeman qubit. We propose a novel protocol for generating two-ion entanglement by collective optical pumping. In our scheme, the optical qubit control and the Zeeman qubit control are combined with controlled dissipation. We experimentally verify the performance of the protocol and produce maximally-entangled states of two ions with 99.3(1)% fidelity. In contrast to the unitary control, we show that collective optical pumping produces the same entangled output state regardless of the input state.

We further explore the toolbox and advantages of multi-level control of trapped ions. We demonstrate the ability to encode, manipulate and readout a qutrit (quantum three-level system) encoded in a single  $^{40}\text{Ca}^+$  ion. This is used to experimentally demonstrate quantum contextuality through a violation of two non-contextuality inequalities.





# Riassunto

Gli ioni atomici intrappolati sono attualmente tra le piattaforme più avanzate per l'elaborazione dell'informazione quantistica. Se da un lato è stato dimostrato che è possibile realizzare tutti i singoli elementi di un computer quantistico universale, aumentare la dimensione del sistema fino renderlo computazionalmente utile rimane ancora una sfida eccezionale.

Il lavoro presentato in questa tesi contribuisce allo sviluppo di processori scalabili per l'informazione quantistica basata su ioni intrappolati. Discutiamo la costruzione e il funzionamento di un sistema a vuoto criogenico (4K) progettato per testare nuove trappole ioniche. Presentiamo il sistema di laser utilizzati per il raffreddamento, la manipolazione e la misurazione degli ioni, con particolare enfasi sullo sviluppo di un laser di controllo del qubit ad alta potenza e ridotta larghezza di riga.

La trappola ionica al centro del sistema criogenico è un nuovo tipo di trappola ad elettrodi di superficie con fotonica integrata. Contrariamente agli esperimenti precedenti, nei quali tutti i fasci laser venivano direttamente allineati sugli ioni, nel nostro apparato la luce di controllo del qubit è inviata attraverso fibre ottiche, accoppiata in guide d'onda integrate nella trappola, e focalizzata sugli ioni usando emettitori dielettrici a reticolo. Questo permette di ridurre le fluttuazioni di puntamento del fascio e di focalizzare fortemente la luce laser, il che si traduce in un alto tasso di accoppiamento laser-ione.

Abbiamo accuratamente caratterizzato diverse trappole con fotonica integrata, trovando numerose sfide associate all'accumulo di cariche nel materiale e al tasso di riscaldamento degli ioni. Tuttavia, abbiamo dimostrato di poter ottenere un alto grado di controllo dello stato quantistico degli ioni utilizzando la luce integrata nella trappola. In particolare, abbiamo implementato porte logiche a due qubit basate sul meccanismo di Mølmer-Sørensen, utilizzate per preparare stati massimamente correlati (entangled) di due ioni  $^{40}\text{Ca}^+$  con fedeltà del 99.45(13)%, limitata da fonti di errore ben caratterizzate.

Oltre al controllo del qubit ottico basato su laser, usiamo campi magnetici a radiofrequenza globali per manipolare il qubit Zeeman. Proponiamo un nuovo protocollo per la generazione di entanglement tra due ioni mediante pompaggio ottico collettivo. Nel nostro schema, il controllo del qubit ottico e il controllo del qubit Zeeman sono combinati con un meccanismo di dissipazione controllata. Abbiamo verificato sperimentalmente le prestazioni del protocollo, e abbiamo preparato stati massimamente correlati di due ioni con fedeltà del 99.3(1)%. Contrariamente al caso di controllo unitario, dimostriamo che il controllo mediante pompaggio ottico collettivo produce lo stesso stato entangled in uscita indipendentemente dallo stato di ingresso.

Abbiamo ulteriormente esplorato le possibilità e i vantaggi del controllo di più livelli

## RIASSUNTO

elettronici, dimostrando la possibilità di codificare, manipolare e leggere un qutrit (sistema quantistico a tre livelli) supportato da un singolo ione  $^{40}\text{Ca}^+$ . Grazie a questo, proponiamo una dimostrazione sperimentale di contestualità quantistica basata sulla violazione di due disuguaglianze di non contestualità.

# Acknowledgements

This work wouldn't have been possible without the help and support of many people.

First and foremost, I'd like to thank Chi Zhang. It's been a privilege building this setup alongside you, and watching you explain the all perplexing mysteries of the experiment one by one with some insane maths. Thanks for keeping the spirits high while everything seems to be falling apart.

I'd like to thank Karan Mehta for the last few years of guidance. Having to convince you that something makes sense before doing it has been both extremely productive and very insightful. I don't think I really know how to "do science" yet, but the parts I know, I have learned from you!

Joseba Alonso – thank you for showing me the ropes of ion trapping, optics, and cryogenics. You and Florian have built an impressive system, and your effort enabled all of these results.

Florian Leupold – thank you for all the supervision, frank conversations, setting an example of an organised scientist with work-life balance, and Mathematica packages that I use until today.

Matt Grau – thank you for \*always\* taking the time to help me with literally anything. It's very clear you'll be a fantastic professor and an amazing father!

Thanks to my office crew – Matt Grau, Vlad Negnevitsky, Thomas Lutz, and Martin Wagener – for all the unproductive work hours. You guys made coming to work really fun, and it was even more fun once the "funbox" was installed.

Thanks for the climbing crew – Christoph Fisher, Martin Wagener, Chris Axline, plus a number of you who enjoy climbing but not to the point of obsession – for many memorable adventures. It was awesome to see climbing grow from "a thing Maciej does that Matteo occasionally joins" to essentially a group sport. Keep it up!

Oliver Wipfli – thank you for all the memorable days working on different experiments in the same room. Hanging out always made me smile. Don't let that cavity break your spirit!

Mirjam Bruttin – thank you for making our work possible, fixing the situation whenever I forgot to pay an invoice, remembering all the important occasions, and keeping the group together!

Thanh Long Nguyen – thank you for your wizardry with construction, electronics and QuTiP. We were very lucky to have you work with us, and apologies for not providing you with long ion strings to run your experiments on in exchange.

Thanks to the students I had a pleasure to supervise: Andrei Militaru, Stefanie Miller, Sebastian Saner, Vilasini Venkatesh, and Kevin Reuer. You guys were all fantastic, and it was really enjoyable to work with you.

## ACKNOWLEDGEMENTS

Shreyans Jain – thank you for your friendship, great conversations, and your bold attitude!

Roland Matt – thanks for making me lough out loud and teaching me physics, often at once.

Thanks in addition to Christa Flühmann, Frieder Lindenfelser, Matteo Marinelli, David Nadlinger, and Robin Oswald. You guys made me feel very welcome when I first joined the group, and I really enjoyed our chats and all the time spent together.

Thanks to all of you who contributed to the work shown in this thesis in ways big and small. Special acknowledgements go to Martin Stadler and Vlad Negnevitsky for their tremendous work on the control system, and patience chasing all the bugs that I introduced. Thank you to my theory collaborators, Florentin Reiter and Ivan Rojko, for helping me understand dissipative processes. Thank you to Daniel Kienzler for all the extremely helpful guidance and advice.

Thank you to Carmelo Mordini and Alfredo Ricci Vasquez. I couldn't have hoped for the experiment to end up in better hands, and I really enjoyed the relatively brief period of working together.

Mirjam Bruttin – thank you for keeping the group running, remembering all the important occasions, and always sorting out my invoicing mistakes! I can't even count how many times you've helped me with issues big and small, and I always enjoy talking with you.

Thanks to all the others who made the TIQI group such a nice place to work – Utku Altunkaya, Gillen Beck, Tanja Behrle, Celeste Carruth, Peter Clements, Jeremy Flannery, David Holzapfel, Luca Huber, Francesco Lancellotti, Hsiang-yu Lo, Brennan MacDonald-de Neeve, Cagri Oenal, Henry Passagem, Simon Ragg, He Ran, Tobias Sägeser, Ilia Sergachev, Nick Schwegler, and Klara Theophilo.

Thank you to my family. Just like you, I know it's not easy to be away from those dearest to you. Thank you for all those moments we get to spend together, and for your support in between.

Last but not least, thank you to ♡ Chiara Decaroli ♡. PhDs have their ups and down, but the last 5 years have been the happiest of my life. I'm very lucky you said "yes", and even luckier that later you said "I do". I'm looking forward to spending the rest of our lives together!

# Contents

<b>Abstract</b>	<b>iii</b>
<b>Riassunto</b>	<b>v</b>
<b>Acknowledgements</b>	<b>vii</b>
<b>List of Acronyms</b>	<b>xv</b>
<b>1. Introduction</b>	<b>1</b>
<b>2. Quantum information with individual trapped ions</b>	<b>7</b>
2.1. Trapping ions . . . . .	7
2.1.1. Why trapped ions? . . . . .	7
2.1.2. Paul trap . . . . .	8
2.1.3. Surface-electrode traps . . . . .	10
2.1.4. Producing ions . . . . .	13
2.2. Controlling ions . . . . .	14
2.2.1. Level structure . . . . .	14
2.2.2. Electric dipole transitions . . . . .	15
2.2.3. Electric quadruple transitions . . . . .	17
2.2.4. Magnetic dipole transitions . . . . .	18
2.2.5. Sideband transitions . . . . .	19
2.2.6. Full Hamiltonian . . . . .	22
2.2.7. Encoding a qubit . . . . .	22
2.2.8. Closed and open-system dynamics . . . . .	23
2.3. Quantum control . . . . .	24
2.3.1. Rabi oscillations . . . . .	25
2.3.2. AC Stark shifts . . . . .	26
2.3.3. Transition broadening and off-resonant excitations . . . . .	28
2.3.4. Decay and saturation . . . . .	29
2.3.5. State preparation . . . . .	30
2.3.6. Optical qubit read-out . . . . .	31
2.3.7. Zeeman qubit read-out . . . . .	34
2.4. Entangling ions . . . . .	35
2.4.1. State-dependent force . . . . .	35

## CONTENTS

2.4.2.	MS gate . . . . .	37
2.4.3.	Phase-modulated Mølmer-Sørensen (MS) gates . . . . .	40
2.4.4.	Measuring entanglement . . . . .	42
2.5.	Cooling ions . . . . .	45
2.5.1.	Role of ion temperature . . . . .	45
2.5.2.	Doppler cooling . . . . .	48
2.5.3.	EIT cooling . . . . .	49
2.5.4.	Resolved sideband cooling . . . . .	51
<b>3.</b>	<b>Experimental setup</b>	<b>53</b>
3.1.	Cryogenic vacuum system . . . . .	53
3.1.1.	Pulse tube stabilisation . . . . .	58
3.1.2.	Helium dewar pressure control . . . . .	60
3.1.3.	Modular radiation shields . . . . .	62
3.1.4.	Two-piece outer vacuum chamber . . . . .	63
3.1.5.	Room-temperature oven . . . . .	63
3.1.6.	Pumping and cool-down . . . . .	65
3.2.	4 K chamber . . . . .	65
3.2.1.	Magnetic field shielding . . . . .	67
3.2.2.	Welded shell . . . . .	68
3.2.3.	Chamber panels . . . . .	70
3.2.4.	Trap mounting . . . . .	71
3.2.5.	Atom sources . . . . .	71
3.3.	Ion traps . . . . .	75
3.3.1.	Gold on quartz trap . . . . .	75
3.3.2.	CMOS trap . . . . .	76
3.3.3.	Waveguide-integrated traps . . . . .	77
3.4.	Electrical control . . . . .	80
3.4.1.	Cryogenic electronics . . . . .	80
3.4.2.	Signal delivery . . . . .	85
3.4.3.	Magnetic field coils and rings . . . . .	90
3.4.4.	Device control . . . . .	92
<b>4.</b>	<b>Free-space and integrated optics</b>	<b>97</b>
4.1.	Basic laser control . . . . .	97
4.1.1.	Gaussian beams . . . . .	97
4.1.2.	Frequency stabilisation . . . . .	100
4.1.3.	Acousto-optic modulation . . . . .	102
4.1.4.	Optical fibers . . . . .	104
4.2.	Photoionization lasers . . . . .	107
4.2.1.	$\lambda = 423$ nm beamline . . . . .	107

4.2.2.	375 nm beamline . . . . .	107
4.2.3.	388 nm beamline . . . . .	108
4.3.	Dipole lasers . . . . .	109
4.3.1.	Repumper beamlines . . . . .	109
4.3.2.	397 nm laser . . . . .	109
4.3.3.	Dipole laser frequency control . . . . .	112
4.4.	Quadrupole laser . . . . .	112
4.4.1.	High-finesse Pound–Drever–Hall (PDH) lock . . . . .	112
4.4.2.	Servo bumps . . . . .	113
4.4.3.	Upgraded quadrupole beamline . . . . .	114
4.4.4.	Fibre noise cancellation . . . . .	114
4.4.5.	Diode injection locking . . . . .	117
4.4.6.	Tapered amplifier . . . . .	122
4.4.7.	Pulsing and multi-tone generation . . . . .	123
4.4.8.	Laser spectrum . . . . .	124
4.5.	Light delivery and collection . . . . .	126
4.5.1.	Free-space light delivery . . . . .	126
4.5.2.	Ion imaging . . . . .	131
4.5.3.	Imaging and laser alignment . . . . .	131
4.5.4.	Ablation laser alignment . . . . .	133
4.6.	Trap-integrated optics . . . . .	134
4.6.1.	Waveguides vs optical fibres . . . . .	135
4.6.2.	Waveguide modes . . . . .	135
4.6.3.	Trap with integrated waveguides . . . . .	137
4.6.4.	Fibre routing . . . . .	142
<b>5.</b>	<b>Single-ion operations with integrated waveguides</b>	<b>145</b>
5.1.	First steps . . . . .	145
5.1.1.	Loading ions . . . . .	145
5.1.2.	Micromotion minimisation . . . . .	146
5.1.3.	Ablation loading . . . . .	147
5.1.4.	Atomic oven loading . . . . .	149
5.2.	State preparation . . . . .	150
5.2.1.	Dipole transition optical pumping . . . . .	151
5.2.2.	Quadrupole transtion optical pumping . . . . .	154
5.3.	State readout . . . . .	156
5.3.1.	Optical qubit readout . . . . .	156
5.3.2.	Zeeman qubit readout . . . . .	158
5.4.	Dephasing and decay . . . . .	159
5.4.1.	Optical qubit lifetime . . . . .	160
5.4.2.	Ramsey experiment . . . . .	161

## CONTENTS

5.4.3.	Optical qubit Ramsey . . . . .	166
5.4.4.	Zeeman qubit Ramsey . . . . .	168
5.5.	Rabi oscillations . . . . .	170
5.5.1.	Optical qubit rotations . . . . .	170
5.5.2.	Rabi frequency noise . . . . .	171
5.5.3.	Zeeman qubit rotations . . . . .	173
5.6.	Cooling and heating . . . . .	175
5.6.1.	Ion thermometry . . . . .	176
5.6.2.	Doppler cooling . . . . .	178
5.6.3.	EIT cooling . . . . .	179
5.6.4.	Heating rates . . . . .	180
<b>6.</b>	<b>Quantum contextuality</b>	<b>185</b>
6.1.	Quantum measurements and entanglement . . . . .	186
6.1.1.	Composite and entangled states . . . . .	186
6.1.2.	Quantum measurements . . . . .	188
6.1.3.	Measurement correlators . . . . .	191
6.2.	Bell theorems . . . . .	194
6.2.1.	Bell Theorem #1 . . . . .	195
6.2.2.	Extended Bell's Theorem and the role of determinism . . . . .	197
6.2.3.	Non-signaling models . . . . .	198
6.2.4.	Nonlocal content and Chained Bell Inequalities . . . . .	200
6.2.5.	Bell inequalities in correlator formalism . . . . .	201
6.3.	Quantum contextuality and KCBS test . . . . .	202
6.3.1.	KCBS inequality . . . . .	204
6.3.2.	Qutrit measurements and quantum violation . . . . .	206
6.3.3.	Exclusivity and connection to Bell inequality . . . . .	209
6.3.4.	Generalised KCBS inequalities and non-contextual content . . . . .	212
6.4.	State-independent contextuality . . . . .	215
6.4.1.	Yu-Oh construction . . . . .	216
6.4.2.	Optimal Yu-Oh inequality . . . . .	219
6.5.	Compatibility loophole . . . . .	220
6.5.1.	The finite precision debate . . . . .	221
6.5.2.	Experimental implications . . . . .	221
6.6.	Is contextuality silly? . . . . .	223
6.6.1.	Toy model of contextual and non-contextual theories . . . . .	224
6.6.2.	When assuming Non-contextuality makes sense . . . . .	224
6.6.3.	When assuming Non-contextuality makes little sense . . . . .	225
6.6.4.	Why perform contextuality experiments? . . . . .	227
6.6.5.	Random number generation? . . . . .	228



<b>7. Experimental studies of quantum contextuality</b>	<b>231</b>
7.1. Ion as a qutrit . . . . .	232
7.1.1. Coherent rotations . . . . .	233
7.1.2. Multi-level readout . . . . .	235
7.1.3. Projective binary readout . . . . .	236
7.1.4. AC Stark shift correction . . . . .	237
7.2. Yu-Oh experiment . . . . .	239
7.2.1. Measurement sequence . . . . .	239
7.2.2. Data acquisition and calibrations . . . . .	241
7.2.3. Purged data . . . . .	241
7.3. Yu-Oh experiment results . . . . .	242
7.3.1. Data analysis . . . . .	242
7.3.2. Assumption checks . . . . .	244
7.4. KCBS experiment . . . . .	245
7.4.1. Measurement construction . . . . .	245
7.4.2. Measurement incompatibility . . . . .	246
7.4.3. KCBS inequality violation . . . . .	249
7.4.4. Comparison with other experiments . . . . .	251
7.5. Generalised KCBS experiment . . . . .	253
7.5.1. Ray construction . . . . .	253
7.5.2. Witness construction . . . . .	254
7.5.3. Experimental results . . . . .	255
<b>8. Entanglement with integrated waveguides</b>	<b>257</b>
8.1. Basic two-ion operations . . . . .	257
8.1.1. Trapping two ions . . . . .	257
8.1.2. Motional modes and cooling . . . . .	259
8.1.3. Optical qubit Rabi oscillations . . . . .	260
8.1.4. Two-ion readout . . . . .	262
8.1.5. Axial stretch mode heating . . . . .	264
8.2. System tuning and MS gate calibration . . . . .	267
8.2.1. Substrate photoconductivity . . . . .	267
8.2.2. Light-induced charging . . . . .	268
8.2.3. Recurrent calibrations . . . . .	271
8.2.4. MS gate calibration . . . . .	274
8.3. MS gates . . . . .	276
8.3.1. Measured performance . . . . .	276
8.3.2. Error model . . . . .	279
8.3.3. Performance in the trap #5 . . . . .	283

<b>9. Optical pumping into a maximally entangled state</b>	<b>287</b>
9.1. Protocol overview . . . . .	289
9.1.1. High-level description . . . . .	290
9.1.2. Theoretical model . . . . .	291
9.1.3. Trapped-ion implementation in $^{40}\text{Ca}^+$ . . . . .	294
9.2. Performance analysis . . . . .	297
9.2.1. Convergence rate . . . . .	297
9.2.2. Effect of errors . . . . .	299
9.2.3. Error simulation . . . . .	304
9.3. System tuning . . . . .	307
9.3.1. Phase-modulated drive . . . . .	308
9.3.2. Spectator modes . . . . .	310
9.3.3. Effective potential anharmonicity . . . . .	313
9.3.4. B-field gradient . . . . .	315
9.3.5. Differential AC Stark shift . . . . .	317
9.3.6. Adapted protocol . . . . .	320
9.4. Experiment and results . . . . .	320
9.4.1. Experimental implementation . . . . .	320
9.4.2. Dynamical behaviour . . . . .	323
9.4.3. Steady-state fidelity . . . . .	323
9.4.4. Understanding errors . . . . .	327
<b>10. Summary and Outlook</b>	<b>331</b>
10.1. Future plans and improvements . . . . .	331
10.1.1. Experimental setup . . . . .	331
10.1.2. Optical setup . . . . .	332
10.1.3. Hybrid qubit operation . . . . .	333
10.1.4. Quantum computing with qudits . . . . .	334
10.1.5. Contextuality tests . . . . .	335
10.1.6. Integrated optics . . . . .	335
10.1.7. Dissipation engineering . . . . .	337
<b>11. Appendix</b>	<b>339</b>
11.1. Clebsch-Gordan coefficients in $^{40}\text{Ca}^+$ . . . . .	339
11.2. Injection lock alignment . . . . .	339
11.3. Undesired FNC back-reflections . . . . .	340
11.4. AOM extinction ratio . . . . .	340
11.5. In-fibre modules for $\lambda = 729\text{ nm}$ laser . . . . .	341
<b>Bibliography</b>	<b>345</b>

# List of Acronyms

AOM	acousto-optic modulator (p. 93, 102–104, 107, 109, 112, 114, 116, 117, 123, 124, 142, 340, 341)
AR	anti-reflection (p. 119, 131)
balun	balanced-to-unbalanced transformer (p. 85)
CW	continuous-wave (p. 111, 117, 123, 134)
DC	static (p. 8, 10, 53, 66, 70, 76–80, 82, 83, 85–89, 95)
DI	device independent (p. 198, 201, 228)
EB	electron-beam (p. 69)
ECDL	external-cavity diode laser (p. 107, 109, 112)
EIT	electromagnetically-induced transparency (p. 49–52, 109, 128, 231, 237)
EOM	electro-optic modulator (p. 101, 337)
FNC	fibre noise cancellation (p. 114, 116, 117, 125, 333, 340, 341)
FWHM	full-width at half-maximum (p. 28, 123, 125)
HV	hidden variable (p. 194, 223, 228)
HWP	half-wave plate (p. 128, 130, 151, 339)
ITO	indium tin oxide (p. 135, 336)
KCBS	Klyachko-Can-Binicioglu-Shumovsky (p. 204, 206–213, 216, 234, 237–239, 245–247, 251, 253, 255)
LED	light-emitting diode (p. 131, 132)
MLE	maximum-likelihood estimation (p. 277)
MS	Mølmer-Sørensen (p. x, 7, 37, 39–42, 44, 93, 104, 114, 257, 309, 314)
NA	numerical aperture (p. 13, 105, 131, 142)

## LIST OF ACRONYMS

NC	non-contextuality (p. 185, 186, 203–206, 208–212, 215, 216, 218, 219, 221–225, 227, 228, 232, 331)
NEP	noise-equivalent power (p. 116)
NS	non-signaling (p. 198–203, 209, 212, 215, 224)
OFE	oxygen-free electronic (p. 68–70, 332)
OSA	optical spectrum analyser (p. 108)
PBS	polarising beam splitter (p. 103, 111, 112, 123)
PCB	printed circuit board (p. 71, 76–78, 82–84)
PCF	photonic-crystal fibre (p. 106–111, 126, 128, 129, 133)
PD	photodetector (p. 101, 113, 114, 116, 117, 122, 123)
PDH	Pound–Drever–Hall (p. xi, 101, 102, 109, 112, 113, 125)
PI	photo-ionisation (p. 109, 126, 128, 129, 133)
PID	proportional-integral-derivative (p. 93, 94, 101, 113, 116, 123)
PM	polarisation-maintaining (p. 106, 107, 109, 111, 123, 126, 128, 130)
PMT	photo-multiplier tube (p. 53, 64, 92, 131)
PR	Popescu-Rohrlich (p. 199)
PRV	pressure relief valve (p. 60, 61)
PT	pulse tube (p. 55–60, 62, 65, 85)
QCCD	quantum charge-coupled device (p. 2, 334)
QEC	quantum error-correction (p. 1, 3, 287)
QKD	quantum key distribution (p. 198, 201)
QM	quantum mechanics (p. 193–195, 197, 198, 200–204, 206, 208–212, 214, 215, 218, 224, 227–229)
QRNG	quantum random number generation (p. 198, 201, 224, 228, 229)
QWP	quarter-wave plate (p. 103, 130, 151, 152, 341)
RF	radio-frequency (p. 8, 10, 53, 65, 70, 76–78, 80, 82, 83, 85, 89, 92–94, 102, 103, 116, 123, 124)

## LIST OF ACRONYMS

RRR	residual-resistance ratio (p. 68, 69)
RT	room-temperature (p. 53, 55)
SET	surface-electrode trap (p. 2, 4, 7, 10–13, 53, 55, 99, 100, 146, 331)
SM	single-mode (p. 105, 107, 129, 135, 137, 138, 142, 341)
SNR	signal-to-noise ratio (p. 333)
SNSPD	superconducting nanowire single-photon detector (p. 135, 336)
SS	stainless steel (p. 59, 60, 63, 64, 71, 74, 75)
TA	tapered amplifier (p. 109, 112, 114, 121–123)
TE	transverse electric (p. 135–137, 139, 140, 142)
TEC	thermoelectric cooler (p. 113)
TM	transverse magnetic (p. 135–137, 140)
TMP	turbo-molecular pump (p. 61, 65)
UHV	ultra-high vacuum (p. 12, 53, 55, 65, 89)
ULE	ultra-low expansion (p. 112)
VCO	voltage-controlled oscillator (p. 114)



# 1. Introduction

A Quantum computer is a device that harnesses the full computational power allowed by the laws of physics to solve otherwise intractable problems. However, it appears that quantum computers are darn hard to build. This thesis describes my contributions to this large and ongoing quest.

The idea of quantum computation was pioneered by P. Benioff, Y. Manin and R. Feynman in the early 1980s [Ben82; Fey82]. In 1985, D. Deutsch made it concrete by proposing the first quantum algorithm, and demonstrating that a quantum computer can do something that a classical computer cannot [Deu85]. Deutsch’s vision<sup>1</sup> sparked a wave of discoveries of possible advantages of quantum computation, such as the Deutsch–Jozsa [Deu92], Bernstein–Vazirani [Ber97], and Simon’s algorithms [Sim97]. The most celebrated results of this era, the quantum search algorithm by L. Grover [Gro96] and the factoring algorithm by P. Shor [Sho94], remain the major driver of a lot of modern quantum computing research.

Back then, it was not obvious whether quantum computing can scale given the presence of unavoidable operation errors [Har96]. This was because, unlike classical information, quantum information cannot be copied [Woo82], and thus classical error correction methods cannot be applied. This changed due to landmark results obtained by P. Shor [Sho95] and independently by A. Steane [Ste96] describing the possibility of quantum error-correction (QEC). Soon it was shown that QEC allows for fault-tolerant quantum computing [Got98; Ste99]. Those results established that quantum computing is feasible in the real world, motivating all subsequent experimental research.

Among the early suggested methods about *how* to perform quantum computation, the 1995 proposal by A. Cirac and P. Zoller for a two-qubit trapped-ion gate was arguably the most influential [Cir95]. Only 7 months after the publication of the proposal, C. Monroe *et al.* published a demonstration of the first two-qubit CNOT gate in a trapped-ion system [Mon95]. By 1997, D. Wineland *et al.* published a review entitled “Experimental issues in coherent quantum-state manipulation of trapped atomic ions” [Win97] (commonly referred to as “The Bible”) which has to this day remained a remarkably complete description of the physics of quantum information processing with trapped ions.

Following seminal demonstrations of multi-ion ground-state cooling [Kin98; Roo99], individual-ion addressing [Näg99] and multi-qubit entanglement [Sac00], trapped-ion research began to focus on how to scale a small quantum information processing unit into a large-scale

---

<sup>1</sup> While Deutsch’s seminal work was motivated by convincing the community of the Many-Worlds interpretation of the quantum theory, that was ironically the one part of his proposal that was not widely accepted.

## 1. INTRODUCTION

processor. The seminal idea put forward by D. Kielpinski *et al.* in [Kie02] was to employ a quantum charge-coupled device (QCCD) architecture, whereby quantum gates are applied locally to small registers of co-trapped ions. Multiple such registers are located on the same trap chip, and ions are physically transported between different registers to connect them. Following the proposal, numerous experiments demonstrated individual components of the QCCD architecture, such as ion splitting and shuttling [Hom09b; Bow12], mixed-species and mixed-isotope cooling [Kie00a; Hom09a], mixed-species gates [Tan15; Hug20], and junction transport [Bla09]. A number of experiments demonstrated small-scale prototypes of a QCCD chip [Hom09b; Pin21].

Large-scale realisation of the QCCD architecture requires scalable trap fabrication and parallelisable optical and electrical control. To facilitate that challenge, J. Chiaverini *et al.* proposed a surface-electrode trap (SET) layout [Chi05], whereby ions are trapped above a planar microfabricated chip. Such traps can be fabricated with high precision in industrial foundries, and accommodate advanced on-chip features. However, SETs are more challenging to operate with low noise, requiring for example cryogenic environments [Lab07] and trap surface treatment [Hit13] to suppress the “anomalous heating” associated with short ion-electrode distances [Tur00]. They are also no panacea, and difficulties scaling up optical and electrical control are among the many reasons the QCCD project has not yet reached its full promise.

This thesis describes our attempts at performing quantum information experiments in cryogenic SETs. Its most important contribution to the quantum computing quest is the first demonstration of two-qubit entanglement driven by laser fields delivered through trap-integrated waveguides [Meh20]. This work, presented in Chapter 8, together with simultaneous complementary demonstration by R. Niffenegger *et al.* [Nif20], shows a path towards scalable optical control of trapped ions. Our work is one of a few examples of functional SETs with integrated features. Recent experiments also demonstrated ion trap operation with chip-integrated voltage sources [Stu18] and chip-integrated detectors [Tod20]. Thus, it appears we are approaching an inflection point where ion trapping may become a fully integrated technology.

It is also very encouraging that despite their limitations, SETs have become commonplace in the community, and we have learned how to manage most of the challenges they impose. Indeed, while until 2019 all the highest-fidelity entangling gates were performed in room-temperature three-dimensional traps [Sac00; Lei03b; Ben08; Bal16; Gae16], recent experiments demonstrated state-of-the-art fidelities in cryogenic SETs with mere 30  $\mu\text{m}$  ion-electrode distance [Sri21; Cla21]. While large traps with tens of ions in a single potential well are still a workhorse of the most advanced quantum simulation experiments [Mon21], the largest quantum volume (a measure of computational power introduced in [Cro19]) reported in any system to date was obtained in a SET [Hon21]. Thus, it is very likely that SETs with advanced on-chip features will lead the way towards large-scale quantum computing with trapped ions.



Trapped ions are only one of the many experimental platforms where quantum computing may be feasible. In the recent years, high-fidelity multi-qubit operations were demonstrated with superconducting qubits [Bar14; She16; Gar20; Kja20], neutral atoms trapped in optical-tweezer arrays [Lev18; Mad20], nitrogen-vacancy centres coupled to nuclear spins in diamond [Ron15; Bra19], and spins in silicon [He19; Hua19]. In photonic platforms, recent works demonstrated a method to construct arbitrarily large two-dimensional cluster states of continuous-variable qubits [Asa19; Lar19], and programmable on-chip quantum information processing within the dual-rail encoding [Qia18]. The project of building a quantum computer is also gradually broadening from universities into startups and industrial labs.

This explosion of interest and competition is valuable, as the field as a whole is nowhere near building a useful universal quantum computer. While state-of-the-art demonstrators operate on few tens of qubits, breakthrough applications in fields such as chemistry will likely require millions [Bab18; vBur20]. This is predominately because even the best quantum gates, with error rates as low as  $10^{-6}$  [Har14], are extremely unreliable when compared to classical gates, whose error rates are as low as  $10^{-17}$  [Har19]. Thus, any algorithmic quantum speedup is reduced by the daunting overhead of QEC, which aims to suppress this noise. For example, according to the resource estimates in [Suc13], for a gate error of  $1 \times 10^{-5}$ , QEC with the surface code increased the qubit count by 3-4 orders of magnitude, and the gate count by 5-6 orders of magnitude compared to an error-free system. As a result, real-world quantum computers may provide no speedup for many problems where ideal quantum computers would [Bab21; Liu21].

The challenge of experimental quantum computing in the year 2021 has therefore two facets. On one hand, the field faces an engineering challenge of implementing the existing ideas with lower noise and at larger scales. At the same time, even 25 years after the Cirac-Zoller proposal, new ideas are needed more than ever. It is unlikely we can “brute-force” our way into a useful quantum-computational advantage. Instead, sustained experimental and theoretical innovation is necessary to decrease the overheads of QEC and overcome some of the long-standing technical roadblocks. For example, recent seminal experiments with trapped ions [Flü18] and superconducting qubits [Gao19] began to explore ways of encoding qubits into quantum harmonic oscillators. This may reduce the overheads of error correction, as a single noisy oscillator may replace multiple noisy qubits [Tzi19]. Another idea, which is explored in Chapter 9, is to engineer dissipation to create entangled states in the steady-state [Ver09]. Dissipative engineering may one day become a highly practical method of creating resource states for quantum error correction [Bag21], and its application in bosonic systems [Kie17; dNee20] is also exciting.

It is worth keeping in mind that, if quantum computing is to fulfil its promise, then 2021 is still just the beginning of the story. Not only is it still unknown what physical system is best suited for quantum computation, but we also understand little about why and when quantum computers excel over classical ones. For example, despite suggestive

## 1. INTRODUCTION

evidence in favour [Joz02; Yog19], it is still unknown whether entanglement is needed for quantum advantage. While Bravyi *et al.* recently proved that there is a complexity separation between some quantum and classical computing models [Bra18; Bra20], the quantum speedup for most problems of practical interest remains a conjecture. Finally, while the “non-classicality” of quantum computing concerns its ability to violate of the extended Church-Turing hypothesis [Aar10], it is still unclear how this relates to the Bell scenario [Bel64] and “quantum contextuality” [Bud21] as a whole, where “non-classicality” refers to the inability ability to reproduce experimental observations through certain hidden-variable models (Chapter 6). In [How14], M. Howard *et al.* linked these notions together by showing that quantum contextuality is a necessary resource for universal quantum computation via magic-state distillation. However, their work also showed it *not* to be a sufficient resource, leaving open the question of what is. Furthermore, to a surprise of many, Zurek *et al.* showed recently that quantum computation can be in fact described through a non-deterministic hidden variable model [Zur20]. These results emphasise that we are yet to grasp where quantum computers get their power from.

This thesis is organised as follows:

**Chapter 2** introduces the basic theoretical concepts of controlling atomic systems. It then discusses the methods of trapping and manipulating individual ions.

**Chapter 3** presents the experimental setup including the cryogenic vacuum system, SETs and signal delivery methods.

**Chapter 4** presents the laser systems used for both free-space and integrated-optical light delivery.

**Chapter 5** presents experimental results trapping, cooling and quantum control of individual  $^{40}\text{Ca}^+$  ions in SETs with integrated waveguides.

**Chapter 6** reviews the theory of quantum non-classicality, with emphasis on quantum contextuality and its experimental testability.

**Chapter 7** presents experimental violations of two non-contextuality inequalities with a single ion.

**Chapter 8** presents experimental results of quantum manipulation of two trapped ions. It demonstrates entanglement generation using laser light delivered through trap-integrated waveguides.

**Chapter 9** introduces and experimentally demonstrates a novel protocol for entangled-state generation using a collective optical pumping process. The experiment uses engineered dissipation to create Bell-states with 93(1)% fidelity, the highest of any dissipative experiment to date.

**Chapter 10** summarises the results and gives an outlook on future work.

I hope you enjoy it!



## 2. Quantum information with individual trapped ions

This chapter presents the basic toolbox of quantum computation with trapped ions. I begin by discussing how ions can be created and trapped, and introduce the specific type of trap we use, called a SET. Sec. 2.2 introduces our ion of choice –  $^{40}\text{Ca}^+$  – and the elements of atomic physics involved in using it. Then I present two ways to encode a qubit that will be used in subsequent chapters. Sec. 2.3 delves into how to perform quantum operations – such as single-qubit gates, state initialisation and readout – on qubits encoded in  $^{40}\text{Ca}^+$ . It also introduces some of the complications that arise due to the presence of spectator atomic transitions. Sec. 2.4 describes the MS gate, which is our two-qubit gate of choice. Finally, in Sec. 2.5 we describe the motivation and methods for cooling the ions to their motional ground states. This chapter serves as a general introduction to subsequent experimental chapters, where some of the results are going to be further refined and compared with observations.

The results discussed in this chapter are also presented and derived in a number of trapped ion review papers and atomic physics textbooks, e.g. [Win97; Lei03a; Häf08; Ste07]. Instead of reproducing all the calculations, I chose to only present the key steps and their implications, and spend more time discussing the intuition behind the equations.

### 2.1. Trapping ions

#### 2.1.1. Why trapped ions?

Quantum systems experience a general trade-off of accessibility versus isolation. On one hand, a qubit should be completely decoupled from its environment in order to maintain coherence. However, a completely isolated system is also completely useless<sup>1</sup>. As soon as a system can be manipulated from the outside, it becomes susceptible to decoherence. Trapped-ion qubits are well-placed on this spectrum: their internal states are extremely well decoupled from the environment, but the use of laser beams propagating through the free space nonetheless allows fast control. As a result, errors are dominated by the imperfections in the control fields. One of the reasons for this isolation is that the trap potential is state-independent, i.e. the ion is trapped in the same well regardless of its electronic state.

---


<sup>1</sup> Furthermore, completely isolated systems probably do not exist in nature. And if they do, it is impossible to find out anything about them!

## 2. QUANTUM INFORMATION WITH INDIVIDUAL TRAPPED IONS

This allows us, for example, to measure the internal state without losing the ion. When multiple ions are trapped together, their internal states do not couple directly, and thus the ion's coherence is not affected by its neighbours. At the same time, high-fidelity qubit-qubit coupling can be engineered on-demand using laser fields at rates of  $> 10$  kHz.

### 2.1.2. Paul trap

The ions in this thesis are confined in an radio-frequency (RF) Paul trap. In a Paul trap, a charged particle is confined by a combination of oscillating (RF) and static (DC) potentials. The linear Paul trap is created by four parallel rods, where an RF voltage is applied to two of them, while the other two are held at RF ground (Fig. 2.1) [Rai92]. This configuration



**Figure 2.1:** Symmetric linear Paul trap. The ion (red) is trapped at the pseudopotential minimum at the centre. The confinement along  $x$  is obtained by applying DC voltages to the endcaps, or to the segmented DC electrode. The DC electrodes are RF grounded but can be offset from the DC ground.

allows one to confine a charged particle in two dimensions ( $y$  and  $z$ ) close to the centre, and by applying a DC voltage to the electrodes in the third dimension  $x$ , the particle can become trapped. The DC voltage can be applied to additional electrodes called the end-caps, or to individual segments of the rod electrodes.

Mathematically we can describe the potential at the centre of the trap as:

$$\Phi(r, t) = \Phi_{\text{DC}}(r) + \Phi_{\text{RF}}(r, t),$$

where the DC potential

$$\Phi_{\text{DC}}(r) = \frac{1}{2} (u_x x^2 + u_y y^2 + u_z z^2)$$

and the RF potential

$$\Phi_{\text{RF}}(r, t) = \frac{1}{2} (v_x x^2 + v_y y^2 + v_z z^2) \cos(\Omega_{\text{RF}} t + \phi).$$

Laplace's equation imposes  $\vec{\nabla}^2 \Phi_{\text{DC}}(r) = \vec{\nabla}^2 \Phi_{\text{RF}}(r, t) = 0$ , implying that  $u_x + u_y + u_z = 0$  and  $v_x + v_y + v_z = 0$ . For an infinitely long symmetric trap, the latter condition implies  $(v_x, v_y, v_z) = v_{\text{RF}}(0, 1, -1)$  for the RF potential. The DC potential is confining along  $x$  (i.e.  $u_x > 0$ ), and one usually sets  $u_y \neq u_z$  in order to break the radial mode degeneracy (Sec. 2.1.3). Given this potential, the ion's time-dependent position  $\vec{r} = (r_x, r_y, r_z)$  can be found by solving the equations of motion

$$\frac{d^2 r_i}{dt^2} = -\frac{e}{m} \frac{\partial \Phi}{\partial r_i}$$

for each direction  $i = (x, y, z)$ . We define

$$a_i = \frac{4eu_i}{m\Omega_{\text{rf}}^2}, \quad q_i = \frac{2ev_i}{m\Omega_{\text{rf}}^2}$$

where  $e$  is the charge and  $m$  is the mass of the ion. Then, the equation of motion can be written as

$$\frac{d^2 r_i}{d\tau^2} + (a_i - 2q_i \cos(2\tau)) r_i = 0 \quad (2.1)$$

where  $\tau = (\Omega_{\text{rf}} t + \pi)/2$ . Eq. (2.1) is a well-studied differential equation known as the *Mathieu equation*, with solutions described e.g. in [McL64]. It turns out that only certain combinations of  $a_i$  and  $q_i$  allow for stable trapping. We work in the limit of  $(|a_i|, q_i^2) \ll 1$ , known as the *first stability regime*. The position of the ion trapped near the trap centre is then well-described by:

$$r_i(t) \propto \cos\left(\frac{\beta_i}{2} \Omega_{\text{rf}} t\right) \left[1 + \frac{q_i}{2} \cos(\Omega_{\text{rf}} t)\right], \quad (2.2)$$

where  $\beta_i = \sqrt{a_i + q_i^2/2}$ . Eq. (2.2) implies that there are two dominant components to the motion of the ion. In the first approximation, the ion is a three-dimensional harmonic oscillator with eigenfrequencies  $\omega_i = \beta_i \Omega_{\text{rf}}/2$ . We can interpret this harmonic oscillation along  $y$  and  $z$  as arising due to a static *pseudopotential* created by time-averaging the RF field. Formally, the RF pseudopotential is given by

$$\Phi_{\text{pp}}(r) = \frac{1}{2} (v'_y y^2 + v'_z z^2)$$

where

$$v'_i = \frac{e}{2m\Omega_{\text{rf}}^2} v_i^2.$$

## 2. QUANTUM INFORMATION WITH INDIVIDUAL TRAPPED IONS

The harmonic oscillator frequencies are then given by

$$\omega_i = \sqrt{\frac{e}{m}} \sqrt{u_i + v'_i}.$$

The oscillation frequency along  $x$  in an infinite-length symmetric linear trap is simply set by the DC curvature, i.e.  $\omega_x = \sqrt{eu_x/m}$ , while the radial eigenfrequencies  $\omega_y$  and  $\omega_r$  are influenced by the RF frequency and curvature, as well as by DC (anti)confinement.

The second component in Eq. (2.2) is the *RF micromotion*, which modulates the slow oscillations at  $\omega_i$  with a component at the RF frequency. For an ion placed at the RF pseudopotential null and cooled down to the motional ground state, this effect can usually be neglected. However, additional micromotion, with sufficient magnitude to significantly modulate the laser-ion interactions, can arise if the DC minimum is not RF null. This can originate from two sources. First of all, more complex trap chips, such as the SETs described in Sec. 2.1.3, do not have a symmetry of the idealised infinite linear trap, and the RF electric field may have a component along the trap axis  $\vec{x}$ . As a result, for a given  $x$ , the pseudopotential minimum may not be a pseudopotential zero. This can also happen in a linear trap due to edge effects of rod misalignment. The micromotion caused by this geometric effect is called the *intrinsic micromotion*<sup>2</sup> [Ami11]. Second, undesired (stray) electric fields with a component along  $\vec{y}$  or  $\vec{z}$  can displace the DC minimum from the RF minimum, causing so-called *excess micromotion*. This effect can be compensated by adding additional static control fields with techniques discussed in Sec. 5.1.2

These calculations also allow us to estimate the amount of confinement produced by an ion trap. For ion-electrode distance  $r_0$  and RF voltage amplitude  $V_{\text{RF}}$ , the RF curvature  $v_{\text{RF}} = d_0 V_{\text{RF}}/r_0$ , where  $d_0 \sim 1$ . The trap depth  $E_{\text{trap}}$  can be estimated by assuming the charged particle needs to reach a ground potential to escape the trap. The associated energy can be calculated as [Pea05]:

$$E_{\text{trap}} = \frac{e^2 V_{\text{RF}}^2}{4mr_0^2 \Omega_{\text{rf}}^2} d_0. \quad (2.3)$$

Typical trap depths of linear four-rod Paul traps are in the order of a few electron-volts. This corresponds to tens of thousands of Kelvin, significantly exceeding the temperatures of atomic ovens and background gas molecules.

### 2.1.3. Surface-electrode traps

In an SET, the tracks defining the RF and DC voltages are all placed on a single planar surface. The DC electrodes are segmented to provide confinement along  $x$ , and more segments allow for finer control, transport, and multizone operation [Wri13; Pin21]. This

---

<sup>2</sup> Some publications, such as [Kel15], use a different nomenclature, where intrinsic micromotion refers to the last term in Eq. (2.2) and is unavoidable even in a perfectly symmetric infinite-length trap. Here, we follow the definition from [Ami11] instead.




geometry was first proposed to enable scalable microfabrication of large traps and trap arrays [Chi05]. In principle, the four rods can be patterned as four tracks on a chip, and the ion can be trapped above. However, in such a design, segmenting the middle DC electrode requires the use of multilayer routing and vias, which is more challenging to manufacture. Instead, experiments typically opt for a five- or six-rod design, where the outer DC lines are segmented, but the inner ones need not be, as shown in Fig. 2.2.



**Figure 2.2:** The electrode arrangement of an SET. The ion position (pseudopotential minimum) is shown in red. The confinement in the plane out of the image is obtained by DC voltages applied to the segmented electrodes. (Left) Five-rod configuration. (Right) Six-rod configuration. Both configuration allow for multi-zone operation without segmenting the middle electrode(s).

As is the case for a four-rod trap, for an infinite-length structure, the pseudopotential is locally symmetric around the trapping position, and the radial modes frequencies are equal unless the DC potential breaks their degeneracy. This is important since a degenerate mode structure cannot be efficiently laser-cooled, see Fig. 2.3 [Ami11]. In a symmetric design,



**Figure 2.3:** The effect of radial mode degeneracy on laser cooling with a single laser parallel to the trap surface. (Left) Radial modes are degenerate. The cooling laser lifts the degeneracy and cools the mode parallel to the trap surface. The mode perpendicular to the trap surface is not cooled and gradually heats up due to spontaneous emission. (Right) Radial modes are non-degenerate and tilted with respect to the trap surface. Now all the modes of motion are coupled to the cooling laser, and the ion can be efficiently cooled close to the ground state.

the degeneracy can be lifted by applying different voltages to opposing DC electrodes. Alternatively, the trap can be made asymmetric, which changes the distances to opposing DC electrodes. The radial eigenaxes should also not be perpendicular to the cooling laser beam (Sec. 2.5.2). In most SET experiments, laser beams travel parallel to the chip surface to avoid scatter. It is then necessary to tilt the eigenaxes away from horizontal/vertical.

Arguably the most challenging aspect of SETs is that the pseudopotential quickly becomes anharmonic as the ion is displaced from the field null. This leads to a significant trap

## 2. QUANTUM INFORMATION WITH INDIVIDUAL TRAPPED IONS

depth reduction, since the energy necessary to move the ion from the RF null to the nearest pseudopotential saddle point is significantly lower than  $E_{\text{trap}}$  in Eq. (2.3) [Pea05]. While it can be optimised, typical trap depths for SETs are a few tens of milli-electron-volts, approximately two orders of magnitude less than linear traps [Hon16]. This has a few important implications. Since the trap depth is now less than the mean energy of the particles effusing from atomic sources, loading SETs is less efficient. Once trapped, the ions are more likely to escape upon collisions with background gases. This effect limits ion lifetimes in typical ultra-high vacuum (UHV) environments to a few hours (with laser cooling constantly on), and cryogenic environments are useful to extend it [Ant09; Lei09].

There are other challenges associated with using SETs for ion trapping. The small gaps between electrodes limit the allowed range of voltages to avoid breakdown [Hug11]. Because of this, high trap frequencies (beneficial to reduce the electric-field noise) require short ion-electrode distances: all of the traps in this thesis were designed for  $h = 50\text{ }\mu\text{m}$  ion height. However, bringing the ion close to the surface increases noise and experimental drifts. This happens for a number of reasons:

1. Electric-field noise increases in the vicinity of the trap surface. Published measurements observed scalings such as  $h^{-2.6}$  [An19],  $h^{-3.5}$  [Des06] and  $h^{-4.0}$  [Sed18b]. This “anomalous heating” can limit the performance of quantum information experiments in SETs. The physical mechanism underlying of this error is not well-understood [Bro14; Ber21], although there is evidence there are multiple mechanisms at play [Sed18b].
2. As  $h$  decreases, it becomes more challenging to align the laser beams onto the ion without scattering light off the chip. Some laser scatter is benign, but blue and ultraviolet photons are found to charge the trap surface in many experiments, as lower-wavelength photons are more likely to exceed the work function of impurities present on the chip electrodes and dielectrics [Wan11; Har10]. This charging changes the electric field and/or its curvature at the ion location, necessitating frequent calibrations [Ami10; Sri21].
3. As  $h$  decreases, a fixed amount of technical noise on the electrodes (e.g. caused by cable pick-up) causes a larger electric field noise, and thus an increased heating rate. Technical noise is not fundamental and can be eliminated through careful design and noise investigation [Sed18a], but still poses an experimental challenge, especially in complex setups with large numbers of connections [Pin21].

All in all, when using SETs, it is simply very easy to stumble upon problems that cause noise, drifts or prevent one from trapping altogether. It wouldn’t all be worth it if they didn’t provide us with serious benefits. Luckily, SETs have a number of strong points acting in their favour:


1. Since they’re open from above, ion fluorescence can be collected with very high

numerical aperture (NA). Some modern SETs include a slot in the middle<sup>3</sup>, allowing for high-NA addressing with low laser scatter as well [Rev20].

2. Short ion-electrode distances mean that high potential curvatures can be produced with modest voltages, making it easier to create dynamic waveforms. It is also indispensable when coupling to ions with near-field microwave gradients [Osp11].
3. By leveraging mature cleanroom fabrication processes, traps can be manufactured at high complexity with low defects, high reliability and yield. This is particularly true for traps fabricated in commercial silicon foundries [Meh14].
4. The fabrication processes allow for integration of potentially game-changing features, such as electronic [Stu18] and optical components [Meh16]. A large fraction of experiments in this thesis were made possible by using integrated optics embedded in the trap.

#### 2.1.4. Producing ions

Every trapped-ion experiment begins with an atomic source. To produce  $^{40}\text{Ca}^+$  ions, we first generate a beam of neutral calcium atoms, either from a resistively-heated atomic oven or from an ablation target [Lei07]. The atoms travelling through the trap region are then ionized with a two-photon process, as illustrated in Fig. 2.4 [Luc03].



**Figure 2.4:** Level structure of neutral calcium. The neutral fluorescence laser at  $\lambda = 423\text{ nm}$  drives a cycling transition, which can be used for isotope selectivity. The second laser at  $\lambda \leq 390\text{ nm}$  ionizes the neutral atom, creating an ion.

A  $\lambda = 423\text{ nm}$  laser drives the  $^1S_0 \leftrightarrow ^1P_1$  transition in neutral calcium. The desired isotope can be selected by tuning the frequency of this laser<sup>4</sup>. The isotope used in this thesis, with the mass number 40, is the most common one, with a natural abundance of 96.9%. Following a state transfer to  $^1P_1$ , a laser with  $\lambda < 390\text{ nm}$  is needed to excite the electron into the continuum<sup>5</sup>. In our experiments, we used both  $\lambda = 375\text{ nm}$  and  $\lambda = 388\text{ nm}$  sources for this purpose, which do not need to be single-mode.

<sup>3</sup> Although users of “real SETs” do not always consider those to be “real SETs”

<sup>4</sup> The shift between subsequent isotopes is  $\approx 300\text{ MHz}$

<sup>5</sup> Jargon for “kicking it out”.

## 2. QUANTUM INFORMATION WITH INDIVIDUAL TRAPPED IONS

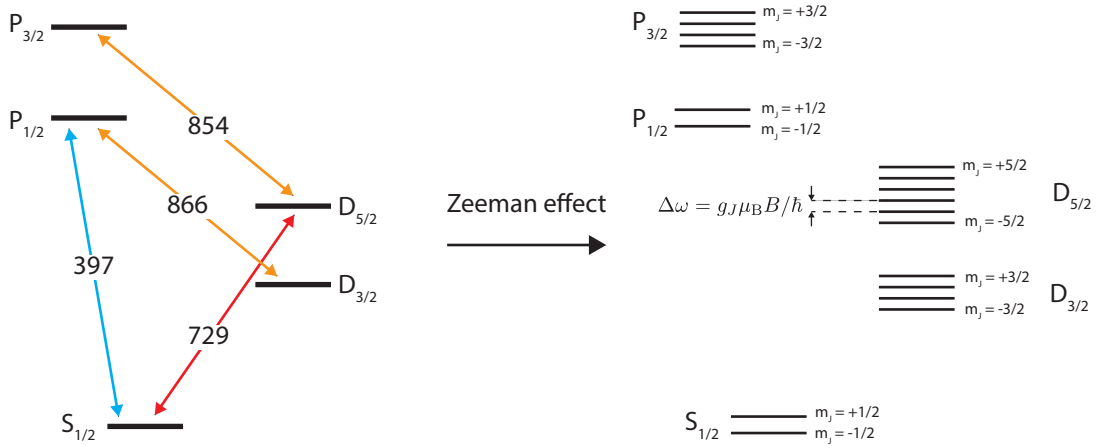
Following photoionization, the ion with sufficiently low initial energy is laser-cooled (Sec. 2.5.2) into the regime where its motion becomes quantum-mechanical. Following a successful loading event, we will see a single particle scattering light (Sec. 2.3.6) into a photon detector. This indicates a presence of an ion. The photoionization lasers are switched off once a desired number of ions is detected.

### 2.2. Controlling ions

Once trapped, the ion's electronic state can be manipulated with electromagnetic fields. We discuss the atomic physics involved in driving internal-state transitions in trapped ions, largely following standard treatments found in [Lei03a] and [Win97].

#### 2.2.1. Level structure

The level structure of a  $^{40}\text{Ca}^+$  ion is shown in Fig. 2.5. The level names are denoted as  $L_J$ , with the total orbital angular momentum  $L$  and the total angular momentum  $J$ . All the levels correspond to different states of a single valence electron, and so the spin quantum number  $S = 1/2$ .



**Figure 2.5:** Calcium ion level structure. (Left) The fine-structure states and the corresponding laser wavelengths (in nm) used in our experiments. The states are denoted by  $L_J$  (Right) The Zeeman effect lifts the degeneracy of the fine structure, splitting each fine-structure state into  $2J + 1$  sub-levels with equal spacing.

We can distinguish two classes of transitions. Lasers at  $\lambda = 393\text{ nm}$ ,  $\lambda = 397\text{ nm}$ ,  $\lambda = 849\text{ nm}$ ,  $\lambda = 866\text{ nm}$  and  $\lambda = 854\text{ nm}$  can excite *electric dipole transitions*. These generally can be driven fast and with low power, but have short lifetimes, see Tab. 2.1 (left). Lasers at  $\lambda = 729\text{ nm}$  and  $\lambda = 732\text{ nm}$  excite *electric quadrupole transitions*. These transitions are forbidden by the dipole selection rules. As a result, the excited states  $D_{3/2}$

and  $D_{5/2}$  are *meta-stable*: they decay to  $S_{1/2}$  is slow, and hence they can be suitable for storing quantum information.

**Table 2.1:** (Left)  $^{40}\text{Ca}^+$  transition wavelengths and linewidths. The top section contains the lasers used in the experiments presented in this thesis (Right) Branching ratios for different decay channels in  $^{40}\text{Ca}^+$ . Within the Zeeman sub-levels, the branching ratios are specified by the Clebsch-Gordan coefficients (App. 11.1)

Transition	$\lambda(\text{nm})$	$\tau(\text{ns})$	$\Gamma(2\pi\text{MHz})$		$S_{1/2}$	$D_{3/2}$	$D_{5/2}$
$S_{1/2} \leftrightarrow P_{1/2}$	397	7.4	21.6				
$S_{1/2} \leftrightarrow D_{1/2}$	729	$1.17 \times 10^9$	$1.36 \times 10^{-7}$				
$P_{3/2} \leftrightarrow D_{5/2}$	854	117.9	1.35	$P_{3/2}$	.92	.0079	.07
$P_{1/2} \leftrightarrow D_{3/2}$	866	107.5	1.48	$P_{1/2}$	.93	.073	
$S_{1/2} \leftrightarrow P_{3/2}$	393	7.4	21.5	$D_{5/2}$	1		
$S_{1/2} \leftrightarrow D_{3/2}$	732	$1.17 \times 10^9$	$1.35 \times 10^{-7}$	$D_{3/2}$	1		
$P_{3/2} \leftrightarrow D_{3/2}$	849	$1.03 \times 10^3$	0.15				

Tab. 2.1 (right) lists the decay channels for different short-lived and meta-stable states. While both  $D_{3/2}$  and  $D_{5/2}$  exhibit a long lifetime, the absence of a  $P_{1/2} \rightarrow D_{5/2}$  decay channel means that the cycling transition at  $\lambda = 397\text{ nm}$  can be used measure whether the ion is in  $D_{5/2}$  or not (Sec 2.3.6). This is the main reason we choose to store quantum information in  $D_{5/2}$ .

In our experiments we use a small static magnetic field  $B$  (between 3.5 G and 6.5 G) to define the quantisation axis, thus lifting the degeneracy of the atomic fine structure. This splits each level  $L_J$  into Zeeman sub-levels of well-defined  $m_J = (-J, -J+1, \dots, +J)$ , as shown in Fig. 2.5. The frequency splitting between adjacent sub-levels is given by:

$$\Delta\omega = g_J \mu_B B / \hbar,$$

where  $\mu_B = 1.4\text{ MHz/G}$  is the Bohr magneton and

$$g_J = \frac{3}{2} - \frac{L(L+1) - S(S+1)}{2J(J+1)}$$

is called the *Lande factor* [Ste07]. The value of  $B$  can be adjusted over a large range depending on experimental desiderata. For very small magnetic fields, it is difficult to spectrally separate different transitions and their sidebands, while for very large magnetic fields, it becomes challenging to address nearby transitions with a single laser beam.

### 2.2.2. Electric dipole transitions

Consider an ion illuminated by a travelling electromagnetic plane wave. We treat the interaction as dominated by the electric field, and ignore the magnetic field component. We also assume the radiation to be classical, which is a good approximation for quasi-classical coherent states of light produced by single-mode lasers. When the light wavelength  $\lambda$

## 2. QUANTUM INFORMATION WITH INDIVIDUAL TRAPPED IONS

significantly exceeds the ion wavepacket size  $z_0$  (we will revisit this assumption later), we can treat the external E-field as a time-dependent perturbation at the ion location:

$$H_1(t) = \frac{eE_0}{2}(r \cdot \epsilon) \cos(\omega t + \phi)$$

where  $r$  is a position operator,  $\epsilon$  is the direction of E-field polarization,  $\omega = 2\pi c/\lambda$  is the plane wave frequency and  $\phi$  is its phase. Consider any two atomic levels  $|\psi_0\rangle, |\psi_1\rangle$  with energies  $E_0$  and  $E_1$  respectively. Their bare Hamiltonian is given by

$$H_0(t) = E_0|\psi_0\rangle\langle\psi_0| + E_1|\psi_1\rangle\langle\psi_1|. \quad (2.4)$$

Let us now expand  $H_1(t)$  in the basis of the two atomic levels. Due to atomic parity,  $\langle\psi_0|H_1|\psi_0\rangle = \langle\psi_1|H_1|\psi_1\rangle = 0$ . Defining  $\tilde{\Omega} = eE_0\langle\psi_0|(r \cdot \epsilon)|\psi_1\rangle/2$ , we can write the perturbation Hamiltonian as

$$H_1(t) = \hbar \left( \tilde{\Omega}|\psi_0\rangle\langle\psi_1| + \tilde{\Omega}^*|\psi_1\rangle\langle\psi_0| \right) \cos(\omega t + \phi). \quad (2.5)$$

We move into the interaction picture with respect to  $H_0$  with a transformation  $|\psi_i\rangle \rightarrow e^{iH_0t}|\psi_i\rangle$ . Defining  $\omega_0 = (E_1 - E_0)/\hbar$  and taking the rotating-wave approximation to eliminate terms oscillating at  $\omega + \omega_0$  gives an interaction Hamiltonian

$$H_I(t) = \hbar \frac{\Omega}{2} |\psi_0\rangle\langle\psi_1| e^{i(\omega - \omega_0)t + \phi} + \text{H.c.}, \quad (2.6)$$

where the dipole Rabi frequency [Jam97]

$$\Omega = \frac{eE_0}{\hbar} |\langle\psi_0|(r \cdot \epsilon)|\psi_1\rangle|. \quad (2.7)$$

and ‘‘H.c.’’ stands for the Hermitian conjugate. Eq. (2.7) describes how the strength of a dipole transition depends on the polarization of the driving radiation and on the properties of the atomic orbitals. For the case of an atom with a single valence electron<sup>6</sup>, transitions with non-zero dipole Rabi frequencies obey the following *dipole selection rules*:

$$\Delta L = \pm 1; \quad \Delta J = 0, \pm 1; \quad \Delta m_j = 0, \pm 1,$$

which explains which transitions in  $^{40}\text{Ca}^+$  are dipole-allowed. Once an atomic matrix element is non-zero,  $\Omega$  will be non-zero if the radiation is polarised along an appropriate direction. By expanding  $r$  in spherical harmonic basis along the quantization axis  $B$  we


---

<sup>6</sup> And hence  $l = L$ , and  $\Delta S = 0$

find that:

$$\begin{aligned} \Delta m_j &= 0 && \text{for } \epsilon \parallel B \\ |\Delta m_j| &= 1 && \text{for } \epsilon \perp B, \text{ circularly polarized.} \end{aligned}$$

Fields which only couple levels with  $\Delta m_j = 0$  and  $|\Delta m_j| = 1$  are known as  $\pi$  and  $\sigma$  drives respectively. Their orientations and polarizations are shown in Fig. 2.6. Any  $\epsilon$  can



**Figure 2.6:** a) Dipole transitions between two levels with  $J = 1/2$ . Blue lines show transitions with  $\Delta m_j = 0$ . b) Beam and polarization ( $\epsilon$ ) orientations which select specific atomic transitions.  $\Delta m_j = 0$  transitions are driven when  $\epsilon \parallel B$ , which requires the beam wavevector to be perpendicular to  $B$ . Transitions with  $\Delta m_j = \pm 1$  are driven by circularly polarized light with  $\epsilon \perp B$ . This requires the wavevector to be along the quantization axis.

be decomposed onto this eigenbasis, which determines the strength with which it drives different transitions.

Finally, we find that the Rabi frequency is closely linked to the transition linewidth  $\Gamma$ . This is unsurprising since the spontaneous emission is a result of a dipolar perturbation driven by the vacuum fields. The difference is that spontaneous emission is amplified by the density of states, which scales as  $\lambda^{-3}$  in free-space. The result is given by the Fermi golden rule

$$\Gamma \propto \frac{\Omega^2}{\lambda^3} \implies \Omega \propto \Gamma^{1/2} \lambda^{3/2}. \quad (2.8)$$

While this is strictly true only for a two-level system, it does offer valuable insights into laser powers necessary to drive different transitions. We calculate these in Sec. 2.3.4.

### 2.2.3. Electric quadrupole transitions

Optical dipole transitions are easy to excite, but the excited state decays very fast, making it a poor contender for storing quantum information. For that purpose, we require long-lived atomic states, transitions between which are likewise harder to drive.

It turns out that many dipole-forbidden transitions can be coupled by the next-order

## 2. QUANTUM INFORMATION WITH INDIVIDUAL TRAPPED IONS

expansion of the electric field. Suppose that the dipole coupling between  $|\psi_0\rangle$  and  $|\psi_1\rangle$  is zero. The second-order perturbation is caused by the instantaneous electric field gradient along the direction of propagation of the plane wave, which is proportional to the wavenumber  $k = 2\pi/\lambda$ . In the rotating-wave approximation, the interaction Hamiltonian is still given by Eq. (2.6), but with a quadrupole Rabi frequency [Jam97]

$$\Omega = \frac{eE_0k}{2\hbar} |\langle\psi_0|(r \cdot \epsilon)(r \cdot k)|\psi_1\rangle|. \quad (2.9)$$

As for the dipole transitions, a number of selection rules can be established for the quadrupole Rabi frequency. These read [Cor06]:

$$\Delta L = \pm 0, 2; \quad \Delta J = 0, \pm 1, \pm 2; \quad \Delta m_j = 0, \pm 1, \pm 2,$$

which explains how  $S_{1/2}$  can be coupled to  $D_{3/2}$  and to  $D_{5/2}$ , and why  $D_{5/2}$  cannot be coupled to  $P_{1/2}$ . The coupling strength is now determined by the direction of both  $\epsilon$  and  $k$  with respect to  $B$ . If  $k \perp B$ , transitions with  $\Delta m = 0$  cannot be driven. If in addition  $\epsilon \parallel B$ , only  $|\Delta m| = 1$  is coupled, while for  $\epsilon \perp B$ , only  $|\Delta m| = 2$  is selected. Transitions with opposite signs of  $\Delta m$  are always coupled with equal Rabi frequency. Full calculation of the geometric factor can be found in [Leu15]. Finally, we note that the Fermi golden rule in Eq. (2.8) applies to quadrupole transitions as well.

### 2.2.4. Magnetic dipole transitions

The Fermi golden rule implies that the transitions with the narrowest linewidth are those with the smallest energy difference. In  $^{40}\text{Ca}^+$  these correspond to different Zeeman sub-levels of the same  $L_J$  manifold. To use these states as a qubit, it is desirable to drive transitions with  $\Delta L = 0$  and at RF/microwave frequencies (MHz to GHz). This is achievable with a pair of optical beams driving a Raman transition [Oze11]. The alternative is to use a magnetic dipole coupling [Kes11]. The magnetic perturbation Hamiltonian is given by

$$H_1(t) = B_0(\mu \cdot \epsilon) \cos \omega t,$$

where  $B_0$  is the strength of the oscillating B-field,  $\mu$  is the magnetic dipole moment of the ion,  $\epsilon$  is the magnetic polarization direction and  $\omega$  is the field oscillation frequency. We find the magnetic dipole Rabi frequency

$$\Omega = \frac{B_0}{\hbar} |\langle\psi_0|(\mu \cdot \epsilon)|\psi_1\rangle|.$$

The appropriate selection rules now read

$$\Delta L = 0; \quad \Delta J = 0, \pm 1; \quad \Delta m_j = 0, \pm 1,$$



allowing to drive transitions within the same Zeeman manifold. Transitions with  $\Delta J = 0$  and  $|\Delta m = \pm 1|$  can be driven most efficiently when  $\epsilon \perp B$ . In such case, the perturbation analysis from Sec. 2.2.2 gives

$$\Omega = \frac{B_0}{\hbar} g_J \mu_B.$$

While the magnetic dipole coupling is caused by every electromagnetic plane wave, it is generally significantly suppressed, since  $B_0 = E_0/c$ . On the other hand, in the low-frequency regime (kHz to GHz), oscillating magnetic fields can be generated by electrical currents, which is particularly efficient in the near field. This allows us to effectively drive magnetic dipole transitions within the same Zeeman manifold.

### 2.2.5. Sideband transitions

We have thus far assumed that the ion is a point particle in free space, and neglected any recoil effects. In reality, the ion is trapped in a three-dimensional harmonic potential, and the interaction with external drive fields affects the motional state as well. Consider a two-level atom coupled to a single harmonic oscillator with frequency  $\omega_m$ . The bare Hamiltonian in Eq. (2.4) gets modified to [Lei03a]:

$$H_0(t) = E_0 |\psi_0\rangle \langle \psi_0| + E_1 |\psi_1\rangle \langle \psi_1| + \hbar \omega_m \left( a^\dagger a + \frac{1}{2} \right). \quad (2.10)$$

The perturbation Hamiltonian  $H_1(t)$  gets modified to account for the direction of travel of the plane wave. The light will perturb the motional state if the wavevector  $\vec{k}$  has an overlap with the oscillation direction. Formally, the ion's position  $\vec{R}$  is given by:

$$\vec{R} = R \hat{r}_0 \quad \text{where} \quad R = z_0 (a^\dagger + a).$$

where  $(a^\dagger + a)$  is the harmonic oscillator position operator,  $\hat{r}_0$  is a unit vector specifying the oscillation direction, and

$$z_0 = \sqrt{\frac{\hbar}{2m\omega}}$$

quantifies the size of the zero-point motion wavepacket ( $z_0 = 11$  nm for  $\omega_m = 1$  MHz). The full perturbation Hamiltonian in Eq. (2.5) gets modified to:

$$H_1(t) = \hbar \left( \tilde{\Omega} |\psi_0\rangle \langle \psi_1| + \tilde{\Omega}^* |\psi_1\rangle \langle \psi_0| \right) \cos(\omega t - \vec{k} \cdot \vec{R} + \phi). \quad (2.11)$$

## 2. QUANTUM INFORMATION WITH INDIVIDUAL TRAPPED IONS

If  $k$  makes an angle  $\theta$  to  $r_0$ , we can re-write the new coupling term for a single motional mode as

$$\vec{k} \cdot \vec{R} = \eta(a^\dagger + a)$$

where the so-called *Lamb-Dicke parameter*

$$\eta = kz_0 \cos(\theta) \quad (2.12)$$

quantifies the strength of the motional coupling ( $\eta \approx 0.07$  for  $z_0 = 11$  nm,  $\lambda = 729$  nm and  $\theta = 45$  deg). Once again we can obtain the full interaction Hamiltonian by moving  $H_1(t)$  in Eq. (2.11) into the interaction picture of  $H_0$  from Eq. (2.10) and taking the rotating wave approximation. We now discover a series of terms which are resonant at different drive frequencies:

$$\begin{aligned} H_I(t) = & \sum_n \frac{1}{2} \hbar \Omega_{n,n} |\psi_0, n\rangle \langle \psi_1, n| e^{i(\omega - \omega_0)t + i\phi} + \text{H.c.} && \text{carrier terms} \\ & + \sum_n \frac{1}{2} \hbar \Omega_{n,n+1} |\psi_0, n+1\rangle \langle \psi_1, n| e^{i(\omega - \omega_0 - \omega_m)t + i\phi} + \text{H.c.} && \text{first blue sideband} \\ & + \sum_n \frac{1}{2} \hbar \Omega_{n+1,n} |\psi_0, n\rangle \langle \psi_1, n+1| e^{i(\omega - \omega_0 + \omega_m)t + i\phi} + \text{H.c.} && \text{first red sideband} \\ & + \sum_n \frac{1}{2} \hbar \Omega_{n,n+2} |\psi_0, n+2\rangle \langle \psi_1, n| e^{i(\omega - \omega_0 - 2\omega_m)t + i\phi} + \text{H.c.} && \text{second blue sideband} \\ & + \dots \end{aligned}$$

where

$$\Omega_{n,m} = \Omega |\langle n | e^{i\eta(a^\dagger + a)} | m \rangle|. \quad (2.13)$$

$\Omega_{n,m}$  can be evaluated analytically for any  $n, m$ , and the resulting Hamiltonian can be simulated once the Hilbert space is truncated. We can however dramatically simplify the problem, and gain intuition in turn, by assuming the ion wavepacket to be much smaller than the wavelength  $\lambda$  of the incident plane wave, which is experimentally well-justified close to the motional ground state. The ion wavepacket extent  $\langle z \rangle = \langle R^2 \rangle^{1/2}$  along  $\hat{r}_0$  can be calculated as an expectation value of the operator

$$z = z_0 \sqrt{(a^\dagger + a)^2}.$$

For example, for a motional Fock state  $n$ , the wavepacket size  $\langle z \rangle = \sqrt{(2n+1)}z_0$ . The motional regime where

$$\langle z \rangle \ll \lambda$$

is known as the *Lamb-Dicke limit*. For those states, the laser-ion interactions can be simplified with an approximation of:

$$\eta\sqrt{(a^\dagger + a)^2} \ll 1.$$

The Rabi frequencies  $\Omega_{n,m}$  in Eq. (2.13) can be then expanded in the first order of  $\eta$ , leading to:

$\Omega_{n,n} = \Omega + O(\eta^2)$	Carrier
$\Omega_{n,n+1} = \Omega\eta\sqrt{n+1} + O(\eta^2)$	Blue sideband
$\Omega_{n+1,n} = \Omega\eta\sqrt{n} + O(\eta^2)$	Red sideband
$\Omega_{n,n+j} = O(\eta^j)$	Higher-order sidebands.

As a result, we can approximate the interaction Hamiltonian with just three terms, each at its own resonant frequency:

$H_I(t) = \hbar\frac{\Omega}{2} \psi_0\rangle\langle\psi_1 e^{i(\omega-\omega_0)t+i\phi} + \text{H.c.}$	Carrier	(2.14)
$+ \hbar\frac{\Omega}{2}\eta \psi_0\rangle\langle\psi_1 a^\dagger e^{i(\omega-\omega_0-\omega_m)t+i\phi} + \text{H.c.}$	Blue sideband	
$+ \hbar\frac{\Omega}{2}\eta \psi_0\rangle\langle\psi_1 ae^{i(\omega-\omega_0+\omega_m)t+i\phi} + \text{H.c.}$	Red sideband.	

Any interaction with a monochromatic drive of frequency  $\omega$  will then be dominated by the closest resonant term, and the other terms will act as spectator modes. A common way to present it is using these three Hamiltonians:

$H_C(t) = \hbar\frac{\Omega}{2}\sigma_+e^{i(-\delta t+\phi)} + \text{H.c.}$	Carrier, $\omega = \omega_0 + \delta$	(2.15)
$H_B(t) = \hbar\frac{\Omega\eta}{2}\sigma_+a^\dagger e^{i(-\delta t+\phi)} + \text{H.c.}$	BSB, $\omega = \omega_0 + \omega_m + \delta$	
$H_R(t) = \hbar\frac{\Omega\eta}{2}\sigma_+ae^{i(\delta t+\phi)} + \text{H.c.}$	RSB, $\omega = \omega_0 - \omega_m + \delta$	

where  $\sigma_+ = |\psi_0\rangle\langle\psi_1|$ . Note that, since the strength of sideband transitions scales as  $\eta$ , it is very inefficient to drive motional couplings with long-wavelength radiation. For example, for a plane wave at the Zeeman qubit splitting frequency of  $\omega = 2\pi \times 16$  MHz and  $z_0 = 11$  nm, the Lamb-Dicke parameter  $\eta < 4 \times 10^{-9}$ . Furthermore, due to the Fock-state dependence of  $\Omega_{n,n+1}$  and  $\Omega_{n,n-1}$ , the mode's temperature can be measured by measuring the motional sidebands. We will present a detailed description of ion thermometry techniques in Sec. 5.6.1, together with experimental results.

### 2.2.6. Full Hamiltonian

Eq. (2.14) is the full approximate Hamiltonian for a two-level system (energy splitting  $\hbar\omega_0$ ) and a single harmonic oscillator (energy splitting  $\hbar\omega_m$ ) interacting with monochromatic radiation. This Hamiltonian is described by three resonances, at frequencies  $\omega_0$  and  $\omega_0 \pm \omega_m$ . To complete the picture, we need to include the effect of more energy levels, multiple ions and more harmonic oscillator states.

If an ion contains  $M + 1$  energy levels  $|\psi_0\rangle, |\psi_1\rangle, \dots, |\psi_M\rangle$  with energies  $E_0, E_1, \dots, E_M$ , then each of the  $M(M - 1)/2$  pairs of states has its own  $H_I(t)$  of the form in Eq. (2.14), and the total interaction Hamiltonian  $H_{I,\text{tot}}(t)$  is the sum of those, with resonances  $\omega_0, \omega_1, \dots, \omega_{M(M-1)/2}$ . The motional state of  $N$  ions is described by a set of  $3N$  harmonic oscillators with frequencies  $\omega_{m,1}, \omega_{m,2}, \dots, \omega_{m,3N}$ . For each ion, the Hamiltonian  $H_I(t)$  is then extended to include  $3N$  blue sideband terms and  $3N$  red sideband terms at frequencies  $\omega_i \pm \omega_{m,j}$ , each with a different Lamb-Dicke parameter. Finally, any single-ion operation on ion  $K$  is constructed by replacing  $|\psi_i\rangle\langle\psi_j| \rightarrow \mathbf{1}^{\otimes K-1} \otimes |\psi_i\rangle\langle\psi_j| \otimes \mathbf{1}^{\otimes N-K}$  in Eq. (2.14), where  $\mathbf{1}$  is the identity operator on the internal states. Multi-ion Hamiltonians are obtained by adding together the constituent single-ion Hamiltonians.

### 2.2.7. Encoding a qubit

In trapped-ion experiments, quantum information can be encoded either in the internal state (electronic levels) or the external state (motional harmonic oscillator) of the ion [Flü18]. In this thesis, we used the former, and the motional state was only used to mediate ion-ion couplings.

In the upcoming chapters, we use two different types of qubits. The optical qubit is encoded within one transition from the  $S_{1/2}$  to the  $D_{5/2}$  manifold. The example choice in Tab. 2.2 is the transition with the minimum B-field sensitivity. The coherence time of the optical qubit is influenced by both the B-field noise and the laser frequency noise. The Zeeman qubit is encoded within the  $S_{1/2}$  manifold [Oze11], and the qubit spacing and coherence is essentially influenced by the B-field alone<sup>7</sup>.

In our experiments, the Zeeman qubit is a superior way of storing quantum information. This is because, in addition to a prolonged lifetime, it is easier to stabilise the magnetic field than the laser frequency. Furthermore, global single-ion operations can be produced fast and with low noise via magnetic dipole transitions (Sec. 2.2.4). However, the optical qubit and its  $\lambda = 729$  nm drive are still necessary to read out the qubit, as we discuss in Sec. 2.3.6.

As discussed in Sec. 2.4, ion-ion coupling is driven by sideband interactions, which are inefficient for long-wavelength radiation. This prohibits us from entangling Zeeman qubits as easily as the optical qubits. There are two standard ways around this problem. Zeeman qubit sideband coupling can be driven by a pair of Raman lasers or by near-field microwave

---

<sup>7</sup> In our experiments, for the Zeeman qubit, the reference oscillator linewidth is negligible compared to other sources of decoherence

**Table 2.2:** Comparison of the typical encodings of the optical qubit and the Zeeman qubit. The hybrid encoding benefits from the long coherence and the single-qubit drive from the Zeeman qubit but mediates entangling and addressing interactions through the optical qubit.

	Optical qubit	Zeeman qubit	Hybrid
Qubit states	$ 0\rangle =  D_{5/2}, m = -1/2\rangle$	$ 0\rangle =  S_{1/2}, m = +1/2\rangle$	$ 0\rangle =  S_{1/2}, m = +1/2\rangle$
	$ 1\rangle =  S_{1/2}, m = -1/2\rangle$	$ 1\rangle =  S_{1/2}, m = -1/2\rangle$	$ 1\rangle =  S_{1/2}, m = -1/2\rangle$
Single-qubit drive	$\lambda = 729$ nm (electric quadrupole)	RF field (magnetic dipole) or Raman transition	RF field (magnetic dipole)
State preparation	Optical pumping	Optical pumping	Optical pumping
Read-out	State-dependent fluorescence	Via optical qubit	Via optical qubit
Lifetime	1.1 s	$\infty$	$\infty$
Coherence	Laser limited	B-field limited	B-field limited
Entanglement	$\lambda = 729$ nm sideband coupling	RF near-field gradient or Raman sideband coupling	Via optical qubit

gradients. These techniques also allow for single-ion addressing. In this thesis we began developing an alternative approach, using a hybrid Zeeman-Optical qubit<sup>8</sup>. This encoding borrows the coherence and single-qubit drive quality from the Zeeman qubit but uses the optical qubit to mediate entangling interactions and single-ion addressing. In Chapter 9 we describe a method of dissipative pumping into entangled states which utilises the hybrid encoding. In the thesis of Chi Zhang, he will describe a hybrid Zeeman-Optical qubit gate, which allows for universal computation in the hybrid qubit encoding [Zha21].

### 2.2.8. Closed and open-system dynamics

For a closed system, the full interaction Hamiltonian  $H_I$  provides all the information necessary to compute the system’s evolution. The time-dependent unitary transformation  $U(t)$  is then given by:

$$U(t) = \exp \left( \int_0^t dt' H_I(t') \right) \quad (2.16)$$

which can be directly integrated if  $H_I(t')$  commutes with itself at different times, or expanded into Magnus series (Eq. (2.26) in Sec. 2.4.1) if it does not. Given  $U(t)$ , an initial pure state

<sup>8</sup> To be clear, the “hybrid qubit” is a set of techniques to perform computation on the Zeeman qubit, not a different qubit per se

## 2. QUANTUM INFORMATION WITH INDIVIDUAL TRAPPED IONS

**Table 2.3:** Lindblad operators for different error channels for the qubit (top rows) and the oscillator (bottom rows) [Tur00].  $\gamma_k$  lists some common names of relevant time-scales. All the noise processes are assumed to be Markovian.

Error name	$L_k$	$\gamma_k$	Example causes
Qubit decay	$\sigma_-$	$1/T_1$	Spontaneous emission
Qubit dephasing	$\sigma_z$	$1/T_2$	Drive frequency noise Magnetic field noise
Motional heating	$\{a_m^\dagger, a_m\}$	$\{\dot{n}, \dot{n}\}$	Electric field noise at $\omega_m$ Spontaneous emission on phonon sideband
Motional dephasing	$a_m^\dagger a_m$	$\Gamma_{\text{dephase}}$	Motional frequency fluctuations

$|\psi(0)\rangle$  at  $t = 0$  evolves into

$$|\psi(t)\rangle = U(t)|\psi(0)\rangle.$$

at time  $t$ .

However, the full interaction Hamiltonian of the system only describes the dynamics that we fully understand and can keep track of. In addition, unwanted environmental interactions induce non-unitary dynamics. These dissipative processes lead to decoherence, which limits the performance of quantum computations. We can simulate the effect of closed or open-system dynamics using a master equation in the Lindblad form [Lin76]:

$$\partial_t \rho = \frac{1}{i\hbar} [H_I, \rho] + \sum_k \gamma_k \left( L_k \rho L_k^\dagger - \frac{1}{2} L_k^\dagger L_k \rho - \frac{1}{2} \rho L_k^\dagger L_k \right). \quad (2.17)$$

Here  $\rho$  is the density matrix of the ion-oscillator system and  $H_I$  is a full interaction Hamiltonian. The non-unitary evolution is described by a set of jump operators  $L_k$  and their associated weights  $\gamma_k$ . For a typical error channel, the decoherence time scale  $t_k = 1/\gamma_k$ . Tab. 2.3 lists the single-ion jump operators corresponding to different common error processes in the Markovian limit. These operators are used in simulations and calculations throughout the thesis.

It should be noted that while dissipation leads to qubit decoherence, it is also a necessary component of any computation. We use dissipation for state preparation (Sec. 2.3.5) and motional cooling (Sec. 2.5). Finally, in Chapter 9, we show how controlled dissipation can be harnessed to directly create entangled states.

### 2.3. Quantum control

We now know how to encode a qubit and how to describe its interaction with external fields. In this section, we describe the toolbox which turns those drives into quantum operations.

### 2.3.1. Rabi oscillations

We begin by showing how the laser-ion interaction Hamiltonian in Eq. (2.6) can be used for single-qubit rotations. Consider two qubits states  $|0\rangle, |1\rangle$  with energy splitting  $\hbar\omega_0$  driven by a laser with frequency  $\omega$  and phase  $\phi$ . Suppose that the detuning  $\delta = \omega - \omega_0$  is significantly below the detuning of any spectator transition, and so we only consider the carrier term from Eq. (2.15):

$$H_I(t) = \hbar \frac{\Omega}{2} \left( \sigma_- e^{-i(\delta t + \phi)} + \sigma_+ e^{i(\delta t + \phi)} \right),$$

where  $\sigma_- = |0\rangle\langle 1|$  and  $\sigma_+ = |1\rangle\langle 0|$ . Using Eq. (2.16), we can derive the associated propagator (in the frame of  $H_0$ ). An alternative method transforms the total Hamiltonian into a time-independent operator. We transform  $H_I(t)$  from the rotating frame of the qubit to the rotating frame of the drive with a unitary  $U(\delta) = |0\rangle\langle 0|e^{i\delta t/2} + |1\rangle\langle 1|e^{-i\delta t/2}$ . The resulting Hamiltonian

$$H'_I = \hbar \frac{\delta}{2} \sigma_z + \hbar \frac{\Omega}{2} \left( \sigma_- e^{-i\phi} + \sigma_+ e^{i\phi} \right) \quad (2.18)$$

is time-independent with  $\sigma_z = |0\rangle\langle 0| - |1\rangle\langle 1|$ . This allows us to write the propagator as  $U = e^{-iH'_I t/\hbar}$ . A simple calculation gives us a propagator for a *Rabi oscillation*:

$$U_I(t) = \begin{pmatrix} \cos(\Omega_\delta t/2) - i(\delta/\Omega_\delta) \sin(\Omega_\delta t/2) & -i(\Omega/\Omega_\delta) \sin(\Omega_\delta t/2) e^{-i\phi} \\ -i(\Omega/\Omega_\delta) \sin(\Omega_\delta t/2) e^{i\phi} & \cos(\Omega_\delta t/2) + i(\delta/\Omega_\delta) \sin(\Omega_\delta t/2) \end{pmatrix} \quad (2.19)$$

where  $\Omega_\delta = \sqrt{\Omega^2 + \delta^2}$  is known as the *generalised Rabi frequency* and we introduced a standard vector notation

$$|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad |1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

In our control system, the phase  $\phi$  is referenced to the drive frequency  $\omega$ , and thus one may look at the drive frame as the natural reference frame for our experimental operations. When  $\delta = 0$ , a qubit in state  $|0\rangle$  undergoes a resonant Rabi oscillation:


$$|\psi(t)\rangle = U_I(t)|0\rangle = \begin{pmatrix} \cos(\Omega t/2) \\ i \sin(\Omega t/2) \end{pmatrix} e^{-i\phi}.$$

Measuring the qubit at time  $t$ , we find the system in  $|0\rangle$  with probability  $P(0)$  given by:

$$P(0) = \cos^2(\Omega t/2) = \frac{1}{2} (1 + \cos(\Omega t)) \quad (2.20)$$

with a full population transfer ( $P(0) = 0$ ) at  $t = \pi/\Omega$ , which is known as the *pi-time*. Example resonant and detuned Rabi oscillations are shown in Fig. 2.7.

## 2. QUANTUM INFORMATION WITH INDIVIDUAL TRAPPED IONS



**Figure 2.7:** Rabi oscillations of a qubit prepared in  $|0\rangle$  at  $t = 0$ . The y-axis shows the probability to measure the system to be at  $|0\rangle$ , given by  $P(0) = |\langle 0 | \psi(t) \rangle|^2$ . Resonant Rabi oscillations (red,  $\delta = 0$ ) are slower and have a large contrast, while off-resonant oscillations (blue,  $\delta = 3\Omega$ ) are faster and have a lower contrast.

In the language of quantum computing, any single-qubit X- or Y rotation can be implemented by adjusting the time and phase of the resonant drive:

$$R(\theta, \phi) = e^{-i\frac{\theta}{2}(\sigma_x \cos \phi + \sigma_y \sin \phi)} = I \cos(\theta/2) + i(\sigma_x \cos \phi - \sigma_y \sin \phi) \sin(\theta/2).$$

with  $\theta = \Omega t$  [Häf08]. We note that single-qubit Z rotation  $Z_\theta = \exp(-i\theta\sigma_z/2)$  generally does not need a physical implementation, as it can be performed by changing the phase of subsequent X and Y rotations (known as a *virtual Z gate*) [McK16].

### 2.3.2. AC Stark shifts

Whenever one transition is driven on resonance, many other transitions are coupled off-resonantly. In the first approximation, those interactions are neglected. However, when considering the limits of control, it is essential to understand the effect of a far off-resonant Rabi drive.

One way to understand single-qubit rotations is through a dressed-state picture. The eigenstates and eigenvalues of  $H'_I$  in Eq. (2.18) are given by:

$$\begin{aligned} |\psi_+\rangle &= \cos(\beta/2) |0\rangle + \sin(\beta/2) |1\rangle & E_+ &= +\hbar\Omega_\delta/2 \\ |\psi_-\rangle &= -\sin(\beta/2) |0\rangle + \cos(\beta/2) |1\rangle & E_- &= -\hbar\Omega_\delta/2 \end{aligned}$$


where the mixing angle  $\beta$  is defined through  $\tan(\beta) = \Omega/\delta$ . On resonance, the eigenstates are split by  $\hbar\Omega$ , which provides an alternative way of understanding Rabi oscillations. In the far off-resonant limit  $\delta \gg \Omega$ , we can expand the solutions up to the leading order in  $\Omega$ ,



resulting in:

$$\begin{aligned}
 |\psi_+\rangle &\approx |0\rangle + \frac{\Omega^2}{2\delta^2}|1\rangle & E_+ &= +\frac{\hbar\delta}{2} + \frac{\Omega^2}{4\delta} \\
 |\psi_-\rangle &\approx |1\rangle - \frac{\Omega^2}{2\delta^2}|0\rangle & E_- &= -\frac{\hbar\delta}{2} - \frac{\Omega^2}{4\delta}.
 \end{aligned} \tag{2.21}$$

The leading-order contribution to the eigenvalues  $\pm\hbar\delta/2$  simply represents an energy shift between the qubit frame and the drive frame, rather than a physical coupling<sup>9</sup>. The next-order term  $\pm\Omega^2/4\delta$  represents a physical interaction known as the *AC Stark shift*. Even in absence of a population exchange between  $|0\rangle$  and  $|1\rangle$ , the qubit states will experience a differential phase shift whenever an off-resonant drive is applied.



**Figure 2.8:** Illustration of the effect of a spectator mode on a qubit transition. When a spectator mode with detuning  $\Delta$  and Rabi frequency  $\Omega_2$  is present in the system, the qubit transition frequency shifts by  $\Omega^2/4\delta$ . This is illustrated as an offset of the dressed state splitting (avoided crossing). The dresses state splitting magnitude is also affected due to mixing with  $|e\rangle$ .

The relevance of the AC Stark shift can be clarified by considering a three-level system shown in Fig. 2.8 (left). Suppose we wish to perform a resonant Rabi oscillation at a Rabi frequency  $\Omega_1$ , which is done by setting  $\delta = 0$  in the two-level picture. However, an additional transition  $|0\rangle \leftrightarrow |e\rangle$  with detuning  $\Delta$  is present in the system. The laser drive couples off-resonantly to that transition with Rabi frequency  $\Omega_2$ , and we assume for simplicity there is no coupling between  $|1\rangle$  and  $|e\rangle$ . In the frame of the drive, the interaction Hamiltonian reads:

$$H_I = \hbar \begin{pmatrix} 0 & \Omega_1/2 & \Omega_2/2 \\ \Omega_1/2 & -\delta & 0 \\ \Omega_2/2 & 0 & \Delta - \delta \end{pmatrix}$$

where we set the zero-point energy to  $E_0$ . By diagonalising  $H_I$  we can find the dressed states in the presence of the spectator mode.

The result is shown in Fig. 2.8 (right). For  $\Omega_2 = 0$ , the dressed-state splitting is  $\hbar\Omega_1$

<sup>9</sup> One way to convince yourself of that is to note that the shift persists in the limit of  $\Omega \rightarrow 0$

## 2. QUANTUM INFORMATION WITH INDIVIDUAL TRAPPED IONS

and the resonance occurs at  $\delta = 0$ . Once  $\Omega_2 > 0$ , the resonance shifts to  $\delta \approx \Omega_2^2/(4\Delta)$ , assuming  $\delta \ll \Delta$ . This means that, to drive the qubit on resonance, we need to set the drive frequency at  $\omega_0 + \Omega_2^2/(4\Delta)$ . Furthermore, if quantum information is encoded in  $|e\rangle$ , driving the qubit on resonance for time  $t$  will cause a phase shift of  $|e\rangle$  of  $\approx \Omega_2^2 t/(4\Delta)$ . This phase offset will be further discussed in Sec. 7.1.4, where it becomes relevant for quantum information processing with three-level systems.

### 2.3.3. Transition broadening and off-resonant excitations

Eq. (2.19) indicates that the contrast of Rabi oscillations is given by  $C = \Omega/\Omega_\delta$ . Expanding this expression for  $\delta \ll \Omega$ , we find that the contrast loss associated with a small detuning is given by

$$1 - C \approx \frac{\delta^2}{2\Omega^2}.$$

This decreased sensitivity to detuning errors with increased Rabi frequency is a result of the *power broadening* of a transition. Increasing  $\Omega$  causes the transition linewidth to increase, which reduces errors associated with detuning (caused by drive or qubit frequency drifts). The power-broadened transition has a full-width at half-maximum (FWHM)  $= 2\Omega$ .

In Sec. 2.3.2 we showed how the dressed state splitting is affected by spectator transitions, but what about the dressed states themselves? Eq. (2.21) implies that a spectator transition with Rabi frequency  $\Omega_2$  and detuning  $\Delta$  mixes into the qubit dressed states with amplitude  $\pm\Omega_2^2/(2\Delta^2)$ . In the computational basis, this corresponds to fast oscillations of the population in  $|e\rangle$  with amplitude  $\sim \Omega_2^2/\Delta^2$ . Consequently, once the pulse is switched off, some of the population might be left in  $|e\rangle$ . This can limit the fidelity of certain interactions, especially when a strong sideband drive off-resonantly excites a spectator carrier mode.

The solution to this problem of *off-resonant mode excitations* is to smoothly shape the amplitudes of the drive pulses. Suppose that a two-level system in  $|0\rangle$  is driven off-resonantly with detuning  $\delta$  and a time-varying Rabi frequency  $\Omega(t)$ . The mixing angle  $\beta(t)$  and the eigenstates  $|\psi_\pm(t)\rangle$  are now time-dependent, and we can write the state at time  $t$  as:

$$|\psi(t)\rangle = A_+(t)|\psi_+(t)\rangle + A_-(t)|\psi_-(t)\rangle$$

By writing out the differential equations for  $A_\pm(t)$ , it can be shown that in the so-called *adiabatic limit*:

$$\left| \frac{d\Omega(t)}{dt} \right| \ll \left| \frac{\Omega_\delta^3(t)}{\delta} \right|$$

the time-dependent coefficients remain constant:

$$A_\pm(t) = A_\pm(0)$$

and thus the system returns to  $|0\rangle$  once the drive is switched off [Sch08]. Of course, no pulse is (nor should be) perfectly adiabatic. Nonetheless, ramping the Rabi frequency at a timescale  $t \sim 1/\delta$  typically does not prolong the interactions significantly, yet significantly suppresses off-resonant excitations [Bal14]. I am not aware of any simple analytical expressions for the residual error magnitude, so I typically simulate it on a case-by-case basis. As an example, the entangling gates described in Sec. 8.3 operate with  $\Omega \approx 2\pi \times 270$  kHz. The gate pulses are near-resonant with the motional sideband which is detuned by  $\delta \approx 2\pi \times 2$  MHz from the carrier. In a square-pulse implementation, this would lead to a significant off-resonant excitation error of  $\delta^2/\Omega_\delta^2 \approx 2\%$ . By smoothly ramping the  $66\text{ }\mu\text{s}$  square pulse over  $2.5\text{ }\mu\text{s}$  we suppress this error by more than three orders of magnitude in the simulation. We note that, in this regime, the exact pulse shape does not matter much, and many window functions can be used to achieve comparable results.

### 2.3.4. Decay and saturation

We have thus far treated the atomic dynamics as a closed system with no decay. We now examine the effect of the finite transition linewidth  $\Gamma$  on Rabi oscillations. This treatment is relevant for both quantum gates (where  $\Gamma \ll \Omega$ ), as well as for driving electric dipole transitions (where  $\Gamma$  and  $\Omega$  are comparable). Formally, we consider a two-level system driven with Rabi frequency  $\Omega$  and detuning  $\delta$ , and solve the master equation with a single jump operator  $L_1 = \sqrt{\Gamma}\sigma_-$  (Tab. 2.3)

In the limit of  $\Gamma \ll \Omega$ , we find that the contrast of the first resonant Rabi flop is reduced to

$$C \approx 1 - \frac{\Gamma}{\Omega}.$$

In Sec. 8.3.2, we will show how this effect sets the lower bound on two-qubit gate errors achievable with a given laser power<sup>10</sup>.

When  $\Gamma$  is comparable or exceeds  $\Omega$ , Rabi oscillations decay fast. Once the coherence is gone ( $t \gg 1/\Gamma$ ), the steady-state is a mixed state with a density matrix

$$\rho = p_0|0\rangle\langle 0| + p_1|1\rangle\langle 1|,$$

where the excited state population [Ste97]:

$$p_0 = \frac{\Omega^2/\Gamma^2}{1 + 2\Omega^2/\Gamma^2}$$

We can calculate the power necessary to efficiently drive a dipole transition by introducing

---

<sup>10</sup> While the limit affects both single- and two-qubit gates, the latter error is bigger by a factor of  $1/\eta$ , and hence more relevant

## 2. QUANTUM INFORMATION WITH INDIVIDUAL TRAPPED IONS

**Table 2.4:** Saturation parameters for dipole transitions in  $^{40}\text{Ca}^+$ . Saturation intensity is calculated from Eq. (2.22), which assumes a random polarisation. Saturation power is calculated at the centre of a Gaussian beam with  $30\text{ }\mu\text{m}$  waist radius.

$\lambda$ (nm)	$I_{\text{sat}}$ (mW/cm <sup>2</sup> )	$P_{\text{sat}}$ (μW)
397	135.4	1.9
854	0.85	0.012
866	0.89	0.013
393	138.9	2.0
849	0.1	0.001

the saturation intensity<sup>11</sup> [Ste07]:

$$I_{\text{sat}} = \frac{\hbar\omega_0^3\Gamma}{4\pi c^2}. \quad (2.22)$$

The ratio  $s_0$  between the drive intensity  $I$  and the saturation intensity  $I_{\text{sat}}$  is known as the *saturation parameter*. It is possible to show that

$$s_0 = \frac{I}{I_{\text{sat}}} = \frac{2\Omega^2}{\Gamma^2},$$

and therefore when  $I = I_{\text{sat}}$  the excited state population  $p_0 = 0.25$ . The saturation intensity can be calculated from the atomic properties (Tab. 2.1), allowing us to infer laser powers necessary to efficiently drive dipole transitions. This is done in Tab. 2.4, where we list both the saturation intensity and the corresponding laser powers for a  $30\text{ }\mu\text{m}$  waist radius.

For a drive detuned by  $\delta$  from the atomic transition, the the steady-state excited population is given by:

$$p_0(\delta) = \frac{1}{2} \frac{s_0}{1 + s_0} \frac{(\Gamma_{\text{eff}}/2)^2}{(\delta^2 + \Gamma_{\text{eff}}/2)^2}$$

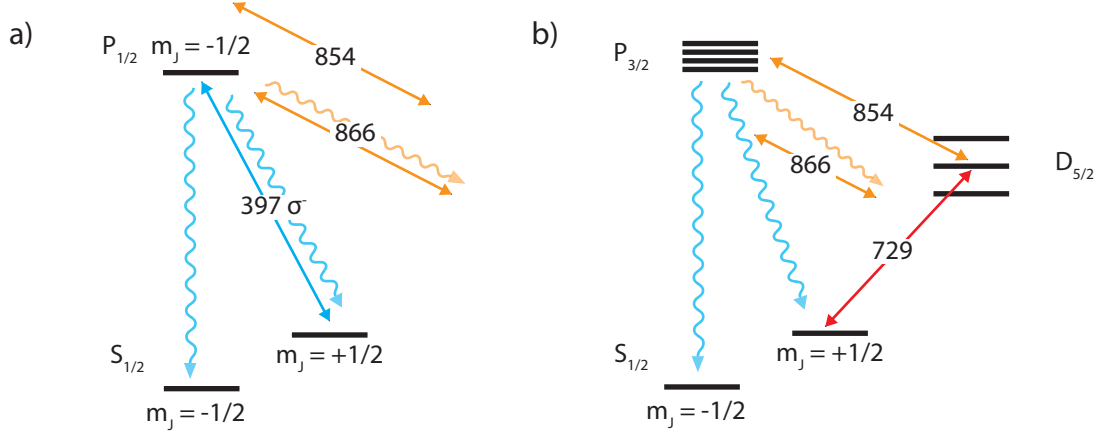
implying that, for a given saturation parameter  $s_0$ , the effective transition linewidth broadens to  $\Gamma_{\text{eff}} = \Gamma\sqrt{1 + s_0}$ . The maximum steady-state population in the excited state is achieved in the limit of  $s_0 \rightarrow \infty$  and  $\delta = 0$ , when  $p_0 \rightarrow 1/2$ .

### 2.3.5. State preparation

The starting step of every quantum computation is to prepare all the qubits in a well-defined initial state. Throughout this thesis, we initialise the ions in  $|1\rangle = |S_{1/2}, m_J = -1/2\rangle$ . The state preparation is the first dissipative protocol we discuss so far.

The first way to initialise the ion is by *optical pumping on a dipole transition* (Fig. 2.9 a) A  $\lambda = 397\text{ nm}$  beam travelling along the quantisation axis is circularly polarised such as

<sup>11</sup> Note that the definition of  $I_{\text{sat}}$  assumes the laser light is unpolarized. If the polarization is aligned for maximum Rabi frequency, the saturation intensity decreases by a factor of 3.



**Figure 2.9:** Different methods of state preparation. a) Optical pumping on a dipole transition uses polarization selectivity to pump into  $|S_{1/2}, m_J = -1/2\rangle$  b) Optical pumping on a quadrupole transition uses frequency selectivity to pump into  $|S_{1/2}, m_J = -1/2\rangle$ .

to only excite the  $|S_{1/2}, m_J = +1/2\rangle \leftrightarrow |P_{1/2}, m_J = -1/2\rangle$  transition. Every spontaneous emission from the excited state populates  $|1\rangle = |S_{1/2}, m_J = -1/2\rangle$  with probability  $1/3$  (this is specified by the Clebsch-Gordan coefficients in App. 11.1). Thus, after approximately three scattering events, the state is prepared in  $|1\rangle$ . This optical pumping is a very robust and reliable method of state initialisation, and we use it throughout the thesis. However, the fidelity is fundamentally limited by the purity of the polarization and by the beam alignment with respect to the magnetic field. We observe this limitation experimentally in Sec. 5.2.1.

These limitations can be overcome by *optical pumping on a quadrupole transition* (Fig. 2.9 b). A  $\lambda = 729 \text{ nm}$  laser beam drives a  $|S_{1/2}, m_J = +1/2\rangle \leftrightarrow |D_{5/2}\rangle$  transition, which is followed by a repump pulse using  $\lambda = 854 \text{ nm}$  and  $\lambda = 866 \text{ nm}$  light. This can be done in a pulsed or continuous fashion, and any  $D_{5/2}$  sub-level except  $|D_{5/2}, m_J = \pm 5/2\rangle$  is suitable<sup>12</sup>. After a few cycles, the ion decays into  $|1\rangle$ . The quadrupole scheme uses frequency rather than polarization sensitivity and hence can achieve much lower errors, limited only by off-resonant excitations. On the other hand, the dipole scheme requires lower powers and is much faster. In experiments, we typically begin with dipole optical pumping, followed by a few cycles of quadrupole pumping to further purify the state (Sec. 5.2).


### 2.3.6. Optical qubit read-out

The optical qubit can be measured by state-dependent fluorescence. Lasers at  $\lambda = 397 \text{ nm}$  and  $\lambda = 866 \text{ nm}$  are applied onto the ion as shown in Fig. 2.10 a). The former drives

<sup>12</sup>  $|D_{5/2}, m_J = +5/2\rangle$  decays with overwhelming probability to  $|S_{1/2}, m_J = +1/2\rangle$  due to dipole selection rules, while the transition from  $|S_{1/2}, m_J = +1/2\rangle$  to  $|D_{5/2}, m_J = -5/2\rangle$  is forbidden by the quadrupole selection rules.

## 2. QUANTUM INFORMATION WITH INDIVIDUAL TRAPPED IONS

continuous excitation and emission on the  $S_{1/2} \leftrightarrow P_{1/2}$  transition<sup>13</sup>, while the latter is applied at high intensity to strongly saturate the  $D_{3/2} \leftrightarrow P_{1/2}$  transition. The fluorescence at  $\lambda = 397$  nm is collected on a photon detector. If the ion started in  $|1\rangle$ , we observe a large number of scattered photons, while the ion in  $|0\rangle$  ideally appears dark [Mye08]. In reality, some photons are also collected in the latter case due to laser scatter and spontaneous emission of  $|0\rangle$ , which is discussed in more detail in Sec. 5.3.1.



**Figure 2.10:** a) Optical qubit readout. Light at  $\lambda = 397$  nm drives a cycling transition, causing the ion to emit fluorescence. Another laser at  $\lambda = 866$  nm repumps the population from  $D_{3/2}$ . b) Zeeman qubit readout can be performed by first applying a number of  $\lambda = 729$  nm shelving pulses, followed by the optical qubit readout.

The measurement is the usually the most efficient when the detection laser at  $\lambda = 397$  nm is applied on resonance ( $\delta = 0$ ) and with power comparable to the saturation intensity. As the drive intensity  $I$  increases above  $I_{\text{sat}}$ , the fluorescence asymptotically approaches the maximum possible value while the background increases linearly with  $I$ . The ideal intensity depends therefore on the amount of measurement background. When the ion is in state  $|1\rangle$ , the mean number of scattered photons measured in a time window  $\tau$  is equal to:

$$\mu_1 = \eta \times \tau \times n_P \times \Gamma \quad (2.23)$$

where  $\eta$  is the overall detection efficiency,  $n_P$  is the fraction of the population in  $P_{1/2}$  and  $\Gamma$  is the transition linewidth. Since the fluorescence does not occur in a pure two-level system, it is actually not possible to obtain  $n_P \rightarrow 0.5$  in the limit of  $I \rightarrow \infty$ . Instead, for the three-level system composed of  $S_{1/2}$ ,  $P_{1/2}$  and  $D_{3/2}$ , we find [Mye08]:

$$n_P = \frac{s_0}{1 + 4s_0},$$

<sup>13</sup> We use a  $\pi$ -polarized beam, but any drive which is not purely  $\sigma^\pm$  suffices

where  $s_0$  is the saturation parameter of the  $S_{1/2} \leftrightarrow P_{1/2}$  transition, and we assumed the  $D_{3/2} \leftrightarrow P_{1/2}$  is maximally saturated. In the limit of  $s_0 \rightarrow \infty$  we find  $n_P = 0.25$ , i.e. the population inversion is half of that available in a two-level system. The intuition about this correction is as follows: when  $s_0 \rightarrow \infty$ , all the sub-levels of the  $S_{1/2}$ ,  $P_{1/2}$  and  $D_{3/2}$  have equal populations. That means half of the population is in  $D_{3/2}$ , which is four-fold degenerate ( $2J + 1 = 4$ ), while  $S_{1/2}$  and  $P_{1/2}$  both have a quarter of the population due to their lower two-fold degeneracy ( $2J + 1 = 2$ ).

Once the fluorescence is collected, we use a simple technique of histogram thresholding to determine the measurement outcome. By preparing and measuring the ion in states  $|1\rangle$  and  $|0\rangle$ , we build a histogram  $H(n)$  of photon counts observed in a time window  $\tau$ . For one ion, we fit the histogram with two Poisson distributions  $H(n) = S(n) + D(n)$ , where

$$D(n) = \frac{\mu_D^n e^{-\mu_D}}{n!}$$

$$S(n) = \frac{\mu_B^n e^{-\mu_B}}{n!}$$

with  $\mu_D$  and  $\mu_B$  as the mean counts for the background and the bright ion respectively. The optimal threshold is then given by [Bur10]:

$$n_{\text{th}} = \frac{\mu_B - \mu_D}{\ln(\mu_B/\mu_D)}.$$

In subsequent experiments, the final state is quantified as  $|1\rangle$  if the number of collected photons exceeds  $n_{\text{th}}$ , and as  $|0\rangle$  otherwise.

By repeating the experiment many times over on the same input state, the thresholding technique allows us to extract the probabilities  $P(0)$  and  $P(1)$  that the ion was in states  $|0\rangle$  and  $|1\rangle$  respectively prior to the measurement. Suppose that the experiment was repeated  $N$  times, with  $n$  trials returning  $|0\rangle$ . Our best estimate of the true probability is the sample mean  $P(0) = n/N$ . The uncertainty can be calculated by noting that each experimental repetition constitutes an independent Bernoulli trial, with the outcomes described by a binomial distribution. Typically, this distribution is well-approximated by a Gaussian distribution (central limit theorem). The resulting uncertainty on  $P_0$  is known as the *quantum projection noise* and is given:

$$\sigma_{P(0)} = \sqrt{\frac{P(0)(1 - P(0))}{N}}. \quad (2.24)$$

The only exception to this rule is when  $P(0) \approx 0$  or  $P(0) \approx 1$ , resulting in a failure of the central limit theorem. In that case, the sample variance calculated from Eq. (2.24) systematically underestimates the population variance for small  $N$ . This is a well-known problem in statistics, and can be solved by applying Laplace's succession rule. This imposes

## 2. QUANTUM INFORMATION WITH INDIVIDUAL TRAPPED IONS

a minimum uncertainty of:

$$\sigma_{P(0),\min} = \frac{1}{N+2}$$

on any probability obtained from an experiment with  $N$  repetitions [Hem14].

Multiple ions in a single crystal can be simultaneously read out by magnifying the image and sending the fluorescence of each ion to a separate photon detector or a different camera pixel. In this thesis, we used a simpler technique of collective fluorescence measurement, where the emission of multiple ions is sent to the same photon counter. For two ions, the detection histogram can be then fitted with three Poisson peaks with means  $\mu_{DD}, \mu_{DB}, \mu_{BB}$ , where  $\mu_{DB} \approx \mu_B$  and  $\mu_{BB} \approx 2 \times \mu_B$ , from which two thresholds are found. By repeating the experiment many times over, we extract the probabilities  $P(00), P(01 + 10)$  and  $P(11)$ . The middle outcome corresponds to the sum of the populations of  $|01\rangle$  and  $|10\rangle$ . While this method can significantly simplify two-ion experiments, it does not scale favourably to larger ion numbers, since a Poisson distribution with mean  $n$  has a variance of  $\sqrt{n}$ . Therefore, for  $N$  ions, while the distribution means are spaced by  $\mu_B$ , their variances increase all the way to  $\sqrt{N \times \mu_B}$  for  $N$  bright ions. For large enough  $N$ , the variance will exceed the peak-to-peak spacing, rendering the readout useless. Already for two ions, we find that the histogram broadening is a limiting factor of our measurement fidelity (Sec. 8.1.4)

The fundamental limit to the readout fidelity comes from non-Poissonian statistics originating from the qubit decay. Finite optical qubit lifetime  $T_1$  introduces a readout error of  $\approx \tau/T_1$  for a detection time  $\tau$ . Consequently, the two-ion peak discrimination error can be reduced by increasing  $\tau$  only until it becomes comparable with the decay error. In Sec. 5.3.1 we quantify the measurement errors introduced by the qubit decay in our experiments.

There are a host of techniques to correct measurement errors. Many of these rely on measurement post-processing – that is, require averaging over multiple runs of the same experiment – while others remove the biases introduced by the thresholding method in a single shot [Mar20]. In this thesis, we only apply a simple linear transformation in post-processing to correct for large readout errors. This method is described in Sec. 5.3 for one-ion readout, and in Sec. 8.1.4 for two-ion readout.

### 2.3.7. Zeeman qubit read-out

Unlike the optical qubit, the Zeeman qubit states cannot be directly distinguished by closed-cycle fluorescence. Instead, we employ auxiliary optical qubit states to perform this detection. We begin the readout with a number of  $\lambda = 729\text{ nm}$  shelving pulses, as shown in Fig. 2.10 b), which transfer the population from  $|0\rangle$  to several states in the  $D_{5/2}$  manifold. Afterwards, we perform the optical qubit measurement as described in the previous section.

While any  $\lambda = 729\text{ nm}$  pulse is prone to errors, shelving can be performed with probability significantly exceeding a single-qubit rotation contrast. This is the case for two reasons.



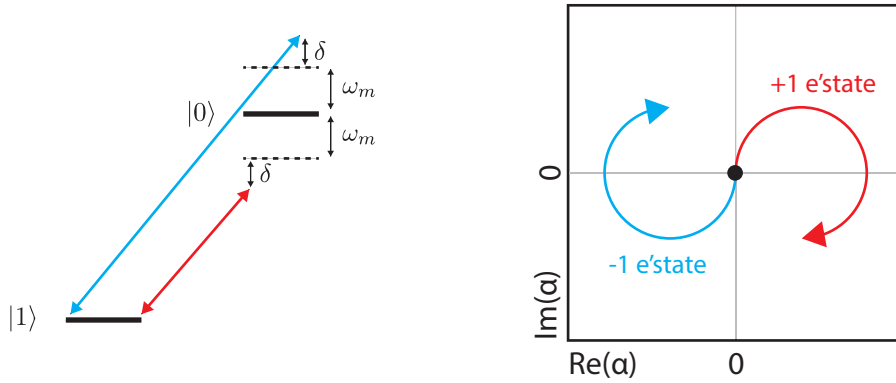
First, since the only goal is to transfer the population, shelving pulses do not require precise phase control. Second, once a pulse was executed, another pulse can be applied on a different transition from  $|1\rangle$  to improve the state transfer probability. Therefore, the Zeeman qubit readout can be usually performed with the same fidelity (but with a slightly longer duration) as the optical qubit readout. Any residual imperfections can also be corrected by post-processing, as discussed in Sec. 5.3.2.

## 2.4. Entangling ions

The ability to entangle ions is a necessary ingredient of all known quantum algorithms. Any pure state quantum computation can be written as an  $N$ -qubit unitary  $U_N$  applied to an initial state  $|0\rangle^{\otimes N}$ , followed by an  $N$ -qubit measurement. The unitary  $U_N$  can be in turn decomposed onto a set of single- and two-qubit operations. In the previous section, we discussed how single-qubit rotations can be implemented on qubits encoded in  $^{40}\text{Ca}^+$  ions. In this section, we describe how to create a two-qubit entangling gate [Bla08]. This gate, together with single-qubit rotations, form a universal set of operations, and any  $U_N$  can be constructed out of them [Häf08].

Of course, entangled states are interesting in their own right. Aside from quantum computation, they are a resource for quantum key distribution [Eke91] and quantum metrology [Tót14]. They are also philosophically important, as they seem to capture the non-classical nature of quantum physics. We discuss the mathematics and philosophy of quantum entanglement in more depth in Chapter 6, while in this section, we focus on the method of applying entangling gates on trapped ions.

### 2.4.1. State-dependent force



**Figure 2.11:** (a) State-dependent force is generated by applying a bichromatic drive with detunings  $\pm(\omega_m + \delta)$  from an optical qubit transition. (b) The resulting trajectory in the position-momentum phase space.

Before discussing the construction of an entangling gate, it is worth understanding the

## 2. QUANTUM INFORMATION WITH INDIVIDUAL TRAPPED IONS

concept of a *state-dependent force* as applied to a single ion. Consider the situation in Fig. 2.11 a), where an ion is simultaneously driven close to the BSB and the RSB. The drives have phases  $\phi_b$  and  $\phi_r$  and detunings  $\pm\delta$  from the BSB and RSB respectively. Ignoring the terms oscillating faster than  $|\delta|$ , the interaction Hamiltonian is given by the sum of two terms from Eq. (2.15):

$$H_I(t) = H_B(t) + H_R(t) = \frac{\hbar\eta\Omega}{2} \left( \sigma_+ a^\dagger e^{-i(\delta t + \phi_b)} + \sigma_+ a e^{i(\delta t - \phi_r)} \right) + \text{H.c.}$$

We define the spin phase  $\phi_s$  and the motional phase  $\phi_m$  as:

$$\phi_s = \frac{1}{2}(\phi_b + \phi_r) \quad \text{and} \quad \phi_m = \frac{1}{2}(\phi_b - \phi_r).$$

This allows us to re-write  $H_I(t)$  as:

$$H_I(t) = \frac{\eta\hbar\Omega}{2} \sigma_\phi \left( a e^{i\delta t} e^{i\phi_m} + a^\dagger e^{-i\delta t} e^{-i\phi_m} \right), \quad (2.25)$$

where  $\sigma_\phi = e^{-i\phi_s} \sigma_+ + e^{i\phi_s} \sigma_-$ . The logic of this substitution is evident from Eq. (2.25): the spin phase enters the spin operator  $\sigma_\phi$  and selects the qubit basis of the interaction, while the motional phase multiplies the oscillator operators only.

We can calculate the effect of applying  $H_I(t)$  using Eq. (2.16). Since the Hamiltonian is time-dependent, we need to be careful regarding operator ordering. One consistent approach is to apply the Magnus expansion [Bla10]:

$$U(t) = e^{-\frac{i}{\hbar} \int_0^t H(t') dt' + \frac{1}{2\hbar^2} \int_0^t \int_0^{t'} [H(t'), H(t'')] dt' dt'' + \dots} \quad (2.26)$$

The first term in the exponential:

$$\begin{aligned} -\frac{i}{\hbar} \int_0^t H(t') dt' &= -\frac{i}{\hbar} \int_0^t \frac{\eta\hbar\Omega}{2} \sigma_\phi \left( \hat{a} e^{i\delta t'} e^{i\phi_m} + \hat{a}^\dagger e^{-i\delta t'} e^{-i\phi_m} \right) dt' \\ &= -\frac{i\eta\Omega}{\delta} \sigma_\phi \sin(\delta t/2) \left( \hat{a} e^{i\delta t/2} e^{i\phi_m} + \hat{a}^\dagger e^{-i\delta t/2} e^{-i\phi_m} \right) \\ &= \left( \alpha(t) \hat{a}^\dagger - \alpha^*(t) \hat{a} \right) \sigma_\phi, \end{aligned} \quad (2.27)$$

where

$$\alpha(t) = -\frac{i\eta\Omega}{\delta} e^{-i\delta t/2} \sin(\delta t/2) e^{i\phi_m}. \quad (2.28)$$

To evaluate the second term, we first calculate the commutator:

$$[H_I(t'), H_I(t'')] = \left( \frac{\eta\hbar\Omega}{2} \right)^2 2i \sin(\delta(t' - t'')). \quad (2.29)$$

where we used the fact that for any  $\sigma_\phi^2 = 1$  for any  $\phi$ . Now we can calculate the second term in the exponential:

$$\begin{aligned} \frac{1}{2\hbar^2} \int_0^t \int_0^{t'} [H(t'), H(t'')] dt' dt'' &= i \left( \frac{\eta\Omega}{2\delta} \right)^2 (\delta t - \sin \delta t) \\ &= i\beta(t), \end{aligned}$$

where I defined

$$\beta(t) = \frac{\eta^2\Omega^2}{4\delta^2} (\delta t - \sin \delta t). \quad (2.30)$$

Finally, observe that the commutator in Eq. (2.29) is a scalar (and not an operator). Therefore, all the further commutators in the Magnus expansion in Fig. (2.26) are zero, and the propagator is exactly given by:

$$\begin{aligned} U(t) &= e^{(\alpha(t)a^\dagger - \alpha^*(t)a)\hat{\sigma}_\phi} e^{-i\beta(t)} \\ &= D(\alpha(t)\sigma_\phi) e^{-i\beta(t)}, \end{aligned} \quad (2.31)$$

where in the last line I introduced  $\hat{D}(\alpha(t)) = e^{(\alpha(t)\hat{a}^\dagger - \alpha^*(t)\hat{a})}$ , a displacement operator which transforms the vacuum state of motion into a coherent state with amplitude  $\alpha(t)$ .

We can now easily interpret the result as shown in Fig. 2.11 b). The initial qubit state can be written as a superposition of the  $\pm 1$  eigenstates of  $\sigma_\phi$ . Under the application of  $H_I(t)$ , these eigenstates become displaced in opposite directions  $\pm\alpha(t)$ , but acquire equal phases  $\beta(t)$ . The displacement occurs in a loop of diameter  $|\alpha| = \eta\Omega/\delta$  in a direction specified by  $\phi_m$ . The displacement returns the motional state to the origin at times specified by:

$$\alpha(t) = 0 \iff t = \frac{2m\pi}{\delta},$$

where  $m \in \mathbb{Z}$ . At those times we find:

$$\beta(t) = \frac{\eta^2\Omega^2 m\pi}{2\delta^2}$$

which is proportional to the area enclosed by the loop in the phase space of position and momentum.

### 2.4.2. MS gate

The MS gate extends the SDF Hamiltonian to two ions addressed with a global beam [Sør99; Sør00]. As discussed in Sec. 2.2.6, we can construct the two-ion Hamiltonian by adding two single-ion Hamiltonians  $H_I(t)$ . This requires a replacement of  $\eta$  with a two-ion Lamb-Dicke

## 2. QUANTUM INFORMATION WITH INDIVIDUAL TRAPPED IONS

parameter

$$\eta = k\sqrt{\frac{\hbar}{2m\omega}} \cos \theta \rightarrow \eta = k\sqrt{\frac{\hbar}{4m\omega}} \cos \theta$$

which reflects the increased total mass<sup>14</sup>. Furthermore,  $\sigma_\phi$  is replaced with a sum of two single-ion operators:

$$\sigma_\phi \rightarrow S_\phi = \sigma_\phi \otimes \mathbb{1} \pm \mathbb{1} \otimes \sigma_\phi.$$

We call  $S_\phi$  the *total spin operator*<sup>15</sup>, and the sign depends on whether the ions oscillate in phase (+ sign) or out of phase (− sign) for the mode in question. All the gates in this thesis were done with the out-of-phase (stretch) axial mode, where  $S_\phi$  has four eigenstates:

$$|\pm\pm\rangle_\phi = \left[ \frac{1}{\sqrt{2}} (|0\rangle \pm e^{i\phi_s} |1\rangle) \right] \otimes \left[ \frac{1}{\sqrt{2}} (|0\rangle \pm e^{i\phi_s} |1\rangle) \right],$$

With eigenvalues:

$$\begin{aligned} S_\phi |+-\rangle_\phi &= 2|+-\rangle \\ S_\phi |-+\rangle_\phi &= -2|-+\rangle \\ S_\phi |++\rangle_\phi &= 0 \\ S_\phi |--\rangle_\phi &= 0. \end{aligned}$$

Aside from this replacement, the first term in the Magnus expansion remains the same. The second term is more complicated since the total spin operator does not square to zero. Instead, we find:

$$S_\phi^2 = 2(\mathbb{1} \otimes \mathbb{1} + \sigma_\phi \otimes \sigma_\phi).$$

Because of this, the commutator is no longer a scalar, but proportional to  $S_\phi^2$ . Still,  $S_\phi^2$  commutes with all other qubit operators in  $H_I(t)$ . Therefore, all the subsequent commutators are still zero, and we can write the total propagator as:


$$U(t) = D(\alpha(t)S_\phi)e^{-i\beta(t)S_\phi^2}.$$

This propagator can generate a maximally entangling gate by setting the detuning  $\delta$  to:

$$\delta = 2\Omega\eta$$

<sup>14</sup> Formally, for any single ion, the Lamb-Dicke parameter is reduced due to a decreased contribution of the motion of any one ion to the normalised normal-mode eigenvector.

<sup>15</sup> Although it is misleading, because the ions are spin-1/2 systems, not spin-1 systems



**Figure 2.12:** Phase space loops corresponding to an MS gate on the stretch mode. a) States  $|+-\rangle$  and  $|+-\rangle$  become displaced, and return to the origin at time  $t = 2\pi/\delta$ , by which time they have accumulated a geometric phase of  $\pi/4$ . The figure shows the displacement right before  $t$ . b) The states  $|++\rangle$  and  $|--\rangle$  remain at the at all times origin, and do not accumulate a geometric phase.

and selecting the pulse time  $t = 2\pi/\delta$  to ensure  $\alpha(t) = 0$ . The resulting propagator then affects the internal state only, and is given by:

$$U(t) = e^{-i\frac{\pi}{8}S_{\phi}^2}.$$

We can understand  $U(t)$  through its action on the internal-state eigenstates, which is illustrated in Fig. 2.12 and evaluates to:

$$\begin{aligned} U(t)|+-\rangle &= e^{-i\frac{\pi}{2}}|+-\rangle = -i|+-\rangle \\ U(t)|-+\rangle &= e^{-i\frac{\pi}{2}}|-+\rangle = -i|-+\rangle \\ U(t)|++\rangle &= |++\rangle \\ U(t)|--\rangle &= |--\rangle. \end{aligned}$$

Such  $U(t)$  implements a controlled-phase gate in the  $|\pm\rangle$  basis, and is a universal gate when combined with single-qubit rotations. We can transform back into the computational basis to find:

$$\begin{aligned} U(t)|00\rangle &= \frac{1}{\sqrt{2}}\left(|00\rangle + ie^{2i\phi_s}|11\rangle\right) \\ U(t)|01\rangle &= \frac{1}{\sqrt{2}}\left(|01\rangle + i|10\rangle\right) \\ U(t)|10\rangle &= \frac{1}{\sqrt{2}}\left(i|01\rangle + |10\rangle\right) \\ U(t)|11\rangle &= \frac{1}{\sqrt{2}}\left(ie^{-2i\phi_s}|00\rangle + |11\rangle\right) \end{aligned}$$

up to a global phase. This representation makes it explicit that  $U(t)$  can transform an initial separable state into a maximally entangled state.

### 2.4.3. Phase-modulated MS gates

One of the main challenges of implementing high-quality entangling gates is in ensuring  $\alpha(t) = 0$  at the end of the sequence. This requires precise calibration of the gate detuning and pulse time and is affected by errors such as trap frequency drifts. Fortunately, the MS gate scheme can be modified to suppress certain errors (usually at a cost of amplifying others).

The first possible modification comes from dividing the single phase-space loop into multiple smaller ones. Formally, we implement a  $K$ -loop MS gate by setting:


$$\begin{aligned}\delta &= \sqrt{K}2\Omega\eta \\ t &= \frac{2\pi K}{\delta}\end{aligned}$$

The resulting unitary is the same as in Eq. 2.4.2. This construction changes the error sensitivity of the gate. Smaller loops suppress the effect of detuning errors and heating rates, but they prolong the gate by a factor of  $\sqrt{K}$ , making it more sensitive to qubit dephasing and uncompensated AC Stark shifts. This is discussed in further in Sec. 8.3.3.

Once the gate is divided into  $K$  loops, another trick can be used to gain further error resilience without losing time. This can be achieved by time-modulation of the drive phases [Hay11]. At any time  $t$ , the direction of the state-dependent force is set by the motional phase  $\phi_m(t)$ , while the spin phase  $\phi_s(t)$  sets the eigenbasis of the interaction. By modulating  $\phi_m(t)$  and keeping  $\phi_s(t)$  constant, the states  $|+-\rangle$  and  $|-+\rangle$  can traverse any loop in the phase space. As long as  $\alpha(t) = 0$  at the end of the interaction, the resulting unitary corresponds to a pure internal-state operator. This allows us to make each loop a different shape, or make each loop run in a different direction.

Multiple schemes have been proposed in the literature to exploit this loop shape freedom to create gates robust to different static or time-dependent errors. In this thesis, we use two different phase modulation methods, each applied to a  $K = 2$ . The first method is the Walsh modulation, which can be considered a kind of spin-echo for displacements [Hay11]. For  $K = 2$ , it involves simply increasing  $\phi_m$  by  $\pi$  between the first and second loop, which changes the direction of the second loop relative to the first. In the presence of timing or detuning errors, this reduces the final residual displacement of  $|+-\rangle$  and  $|-+\rangle$  relative to each other, as illustrated in Fig 2.13.

Following the calculation in Eq. (2.27), a MS gate segment of length running from  $t = 0$



**Figure 2.13:** a) A two-loop MS gate with a fixed detuning error. The first loop does not return to the origin, and the second loop increases the final error. b) A two-loop Walsch-modulated MS gate with a fixed detuning error. The first loop does not return to the origin, at which point the direction of the state-dependent force is reversed. The second loop cancels the error of the first loop in the first order.

to  $t = \tau$  with detuning  $\delta$  and motional phase  $\phi_m$  introduces a displacement given by

$$\alpha(\tau) = k \int_0^\tau e^{i\delta t} e^{i\phi_m} dt$$

where  $k = -i\eta\Omega/2$ . Consider a two-loop modulated gate of length  $t = 2\tau$  with segment length  $\tau = 2\pi/\delta$ . The motional phases are given by

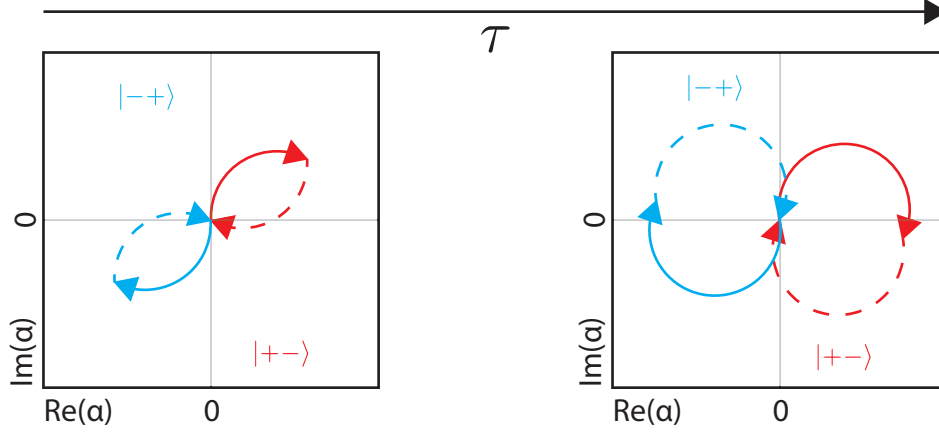
$$\phi_1 = \phi_m \quad \text{and} \quad \phi_2 = \phi_1 + \phi_0$$

for the first and second segment respectively. The final displacement evaluates to:

$$\begin{aligned} \alpha(2\tau) &= k \int_0^\tau dt e^{i\delta t} e^{i\phi_m} + k \int_\tau^{2\tau} dt e^{i\delta t} e^{i\phi_m} e^{i\phi_0} \\ &= -\frac{ie^{i\phi}}{\delta} \left( -1 + e^{i\delta\tau} \right) \left( 1 + e^{i(\delta\tau + \phi_0)} \right) \end{aligned} \quad (2.32)$$

To compare the effect of a detuning error  $\epsilon$  on modulated and unmodulated gates, we can expand this expression around  $\delta = \delta_0 + \epsilon$  for a segment length  $\tau = 2\pi/\delta_0$ . This gives a residual displacement magnitude of

$$\begin{aligned} |\alpha(2\tau)| &= k \frac{4\pi}{\delta_0^2} \epsilon + O(\epsilon^2) & \phi_0 = 0, \text{ unmodulated gate} \\ |\alpha(2\tau)| &= k \frac{4\pi^2}{\delta_0^3} \epsilon^2 + O(\epsilon^3) & \phi_0 = \pi, \text{ modulated gate.} \end{aligned}$$



**Figure 2.14:** Aussie-style phase modulated MS interaction. By adjusting the phase modulation in proportion to the interaction time, the population always returns to the origin.

In other words, adding phase modulation removes the residual spin-motion entanglement to the first order in  $\epsilon$ . More generally, by using  $K$  loops and  $(K - 1)$ –order Walsh modulation, we can reduce this error to  $O(\epsilon^K)$ . This is important, as residual spin-motion entanglement creates a gate error  $\propto |\alpha(2\tau)|^2$  (Sec. 2.5.1).

Sometimes it is valuable to ensure that  $\alpha(t) = 0$  even if the pulse area does not result in a maximally entangling gate. This can be achieved by “Aussie-style” phase modulation [Mil18b]. For  $K = 2$  loops, loop closure is achieved by setting

$$\phi_0 = \pi - \delta\tau \quad \text{modulated gate, Aussie style.}$$

Plugging this to Eq. (2.32) we find that  $\alpha(2\tau) = 0$  for all  $\tau$ , which is illustrated in Fig 2.14. This method of phase modulation allows us to implement a general unitary  $U(t) = e^{-i\beta(t)S_\phi^2}$  where  $\beta(t)$  can be adjusted by simply changing the total interaction time. Sec. 9.3.1 demonstrates the experimental implementation of such phase modulation. Aussie-style phase modulation reduces to Walsh modulation when  $\tau = 2\pi/\delta$ , and provides the same suppression of detuning errors in the leading order.

#### 2.4.4. Measuring entanglement

Since two-qubit gates are typically the most challenging aspects of quantum computing experiments, it is important to measure and benchmark their quality. One method of achieving that is through quantum process tomography, where the full operator description of the process is estimated and compared with the ideal entangling operation. Unfortunately, the full process tomography is challenging to implement for a number of reasons. First, process tomography requires single-qubit rotations to prepare and read out in a tomographically complete basis. This creates more stringent setup requirements than the MS gate itself, which only needs a global beam. While late in the thesis we developed a method of single-ion



addressing, it is currently much noisier than the entangling interaction itself (Sec. 9.3.5). Second, process tomography requires single-ion readout, whereas our experimental setup only measures the total number of ions in the state  $|1\rangle$  (Sec. 2.3.6).

A common workaround in our field is to use a two-qubit gate to prepare a maximally entangled state. The fidelity of the resulting Bell state can be measured with only global rotations and read-out, implementing so-called *partial state tomography*. We then use this fidelity to benchmark the quality of the entangling interaction itself<sup>16</sup>.

Formally, we define the state fidelity as

$$\mathcal{F}(\rho, \phi) \equiv \langle \phi | \rho | \phi \rangle$$

where  $|\phi\rangle$  is the desired pure state and  $\rho$  is the actual density matrix. How can we estimate  $\mathcal{F}$  with only global rotations and measurements? Let us begin by discussing fidelity estimation in the traditional Bell state basis:

$$\begin{aligned} |\Psi^\pm\rangle &= \frac{1}{\sqrt{2}} (|00\rangle + |11\rangle) \\ |\Phi^\pm\rangle &= \frac{1}{\sqrt{2}} (|01\rangle + |10\rangle) \end{aligned} \tag{2.33}$$

This basis can be defined as eigenstates of two-ion parity operators  $\sigma_i \otimes \sigma_i$  for  $i = (x, y, z)$ . A nice property of the parity operators is that they can be measured with only global rotations and readout. Specifically, by post-processing a global readout we extract

$$\langle \sigma_z \sigma_z \rangle = P(00) + P(11) - P(01 + 10)$$

as defined in Sec. 2.3.6. We can likewise extract  $\langle \sigma_x \sigma_x \rangle$  and  $\langle \sigma_y \sigma_y \rangle$  by applying a global  $\pi/2$  “analysis” pulse before the measurement, and then performing a parity readout. Specifically,  $R(\pi/2, \pi/2) \otimes R(\pi/2, \pi/2)$  maps the X-basis onto the computational basis, and allows us to measure  $\langle \sigma_x \sigma_x \rangle$ , while  $R(\pi/2, 0) \otimes R(\pi/2, 0)$  maps the Y-basis onto the computational basis, and allows us to measure  $\langle \sigma_y \sigma_y \rangle$ .

Armed with the ability to measure parities, we can extract the fidelities in the Bell basis. Consider a general two-ion density matrix [Hom06]

$$\rho = \begin{pmatrix} w & p & q & r \\ p^* & x & s & t \\ q^* & s^* & y & u \\ r^* & y^* & u^* & z \end{pmatrix}$$

<sup>16</sup> While it is common for trapped-ion papers to measure the Bell state fidelity and call the outcome “gate fidelity”, it is an unwarranted leap unless the noise model is extremely well understood.

## 2. QUANTUM INFORMATION WITH INDIVIDUAL TRAPPED IONS

where  $w = 1 - x - y - z$ . By direct calculation we find

$$\begin{aligned}\langle \sigma_x \sigma_x \rangle &= r^* + s^* + r + s, \\ \langle \sigma_y \sigma_y \rangle &= -r^* + s^* - r + s, \\ \langle \sigma_z \sigma_z \rangle &= 1 - 2x - 2y.\end{aligned}$$

Similarly, the fidelities of interest can be written as

$$\begin{aligned}\mathcal{F}(\rho, \Phi^\pm) &= \frac{1}{2} (1 - x - y \pm r \pm r^*) \\ \mathcal{F}(\rho, \Psi^\pm) &= \frac{1}{2} (x + y \pm s \pm s^*)\end{aligned}$$

which results in experimentally useful formulae:

$$\begin{aligned}\mathcal{F}(\rho, \Phi^\pm) &= \frac{1}{4} \left[ 1 + \langle \sigma_z \sigma_z \rangle \pm (\langle \sigma_x \sigma_x \rangle - \langle \sigma_y \sigma_y \rangle) \right] \\ \mathcal{F}(\rho, \Psi^\pm) &= \frac{1}{4} \left[ 1 - \langle \sigma_z \sigma_z \rangle \pm (\langle \sigma_x \sigma_x \rangle + \langle \sigma_y \sigma_y \rangle) \right],\end{aligned}\tag{2.34}$$

which imply that three global parity measurements with appropriate settings suffice to determine the Bell state fidelity of any two-qubit state. We will use those equations directly in Chapter 9 to determine the fidelity of  $|\Psi^-\rangle$  obtained by collective optical pumping into the entangled states of the Zeeman qubit.

In Chapter. 8, we benchmark an optical qubit MS gate acting on  $|11\rangle$ . Such interaction does not prepare a Bell state of the form in Eq. (2.33). However, our method can be readily adapted to any maximally entangled state of the form  $|\Phi^\phi\rangle = \frac{1}{\sqrt{2}}(|00\rangle + e^{i\phi_0}|11\rangle)$ . This is because  $|\Phi^{\phi_0}\rangle$  can be transformed into  $|\Phi^\pm\rangle$  by a global Z rotation. As mentioned in Sec. 2.3.1, this rotation can be performed virtually by changing the phase of the analysis pulse. Formally, we define a general measurement basis as  $\sigma_\phi = \cos(\phi)\sigma_y + \sin(\phi)\sigma_x$ . We can map onto this basis with a global  $R(\pi/2, \phi/2) \otimes R(\pi/2, \phi/2)$  pulse. Then, the fidelity between  $\rho$  and a general Bell state  $|\Phi^{\phi_0}\rangle$  is given by:

$$\mathcal{F}(\rho, \Phi^{\phi_0}) = \frac{1}{4} \left[ 1 + \langle \sigma_z \sigma_z \rangle + \langle \sigma_{\phi_0} \sigma_{\phi_0} \rangle + \langle \sigma_{\phi_0+\pi/2} \sigma_{\phi_0+\pi/2} \rangle \right].$$

Note that, while three measurement settings suffice to extract the fidelity, these must be tailored to the expected value of  $\phi_0$ . It may be challenging to determine this phase *a priori*, since it is influenced by the time-dependent AC Stark shifts. Therefore, following a measurement of  $\langle \sigma_z \sigma_z \rangle$ , it is common to simply scan  $\phi$  and find the maximum and minimum values of the parity curve. Thus, the expression for the fidelity is often recast in the form:

$$\mathcal{F}(\rho, \Phi^{\phi_0}) = \frac{1}{2} \left[ P(00) + P(11) + C \right]\tag{2.35}$$

where  $C$  is the contrast of the parity curve:

$$\langle \sigma_\phi \sigma_\phi \rangle = C \cos(2\phi - \phi_0). \quad (2.36)$$

However, for more complicated experiments, it is useful to keep track of  $\phi_0$ . This also allows one to optimise the entanglement fidelity more efficiently.

Finally, can we adapt our global rotations to calculate the fidelity with respect to a general odd Bell state  $|\Psi^\phi\rangle = \frac{1}{\sqrt{2}}(|01\rangle + e^{i\phi}|10\rangle)$ ? Alas, the answer is no. As will be discussed in Chapter 9, global rotations leave  $|\Psi^-\rangle$  invariant, and so we cannot transform to the desired measurement basis at will.

## 2.5. Cooling ions

Most quantum information experiments with trapped ions begin by cooling the ions close to the motional ground state. This is not necessary for trapping – in fact, ions can remain confined for a long time with all the cooling lasers switched off – but it turns out that ion temperature affects the quality of coherent operations. In this section, we review the cooling requirements and show how motional mode cooling is implemented.

### 2.5.1. Role of ion temperature

There are a number of coherent control issues that arise when the ion is far from the motional ground state. The exact amount of cooling required depends on the operation to be implemented and on the desired fidelity. Here, we present an overview of ways in which temperature may induce undesired operation errors.

#### Higher-order sidebands

In writing the approximate light-ion interaction Hamiltonian in Eq. (2.14), we assumed  $\eta\sqrt{(2n+1)} \ll 1$  for all motional modes, which is only the case close to the motional ground state. When the ion is hotter, more terms need to be included in the Hamiltonian, corresponding to higher-order motional sidebands. Once more motional modes are included, this also gives rise to cross-coupling terms between different harmonic oscillators. Since sideband couplings become stronger as  $n$  increases, the problem of spectator transitions becomes accentuated for higher  $n$ . The resulting spectral crowding and spectator mode broadening makes it more challenging to avoid off-resonant excitations and to keep track of the AC Stark shifts.

### Rabi frequency fluctuations

In the first approximation, the carrier Rabi frequency  $\Omega_{n,n}$  in Eq. (2.13) is independent of the motional state  $n$ . Adding the next-order correction gives:

$$\Omega_{n,n} = \Omega \left( 1 - \frac{\eta^2}{2}(2n+1) \right) + O(\eta^4).$$

This correction would not introduce a decoherence path if the ion was in a Fock state. However, since the oscillator state after laser cooling is a thermal distribution, this leads to a shot-to-shot fluctuation in the Rabi frequency. Suppose that the ion is coupled to a single harmonic oscillator mode in a thermal state with average phonon number  $\bar{n}$ . Every shot of the experiment, the mode occupation  $n$  is then sampled from a thermal distribution:

$$P_{\text{th}}(n) = \frac{(\bar{n})^n}{(1 + \bar{n})^{n+1}}. \quad (2.37)$$

This means that the Rabi frequency is sampled from a distribution of

$$\Omega_{\text{th}}(n) = P_{\text{th}}(n) \Omega |\langle n | e^{i\eta(a^\dagger + a)} | n \rangle|$$

We can find the mean Rabi frequency  $\langle \Omega_{\bar{n}} \rangle$  and its variance  $\sigma_{\Omega_{\bar{n}}}$  by evaluating the first and second moments of this distribution. The result is:

$$\begin{aligned} \langle \Omega_{\bar{n}} \rangle &= \Omega e^{-\eta^2(\bar{n}+1/2)} \\ \sigma_{\Omega_{\bar{n}}} &\approx \langle \Omega_{\bar{n}} \rangle \eta^2 \sqrt{\bar{n}(\bar{n}+1)} \end{aligned} \quad (2.38)$$

where the last equation assumes the Lamb-Dicke limit [Win97]. While Eq. (2.38) describes the contribution of a single motional mode, in reality, it must be multiplied over all motional spectator modes of the ions. This result (with small modifications) applies to sideband operations as well. This expression tells us that, as temperatures increase, the shot-to-shot Rabi frequency fluctuations increase as well. Therefore, fidelity of quantum operations can be increased by ground-state cooling spectator modes with significant Lamb-Dicke parameters.

### Increased entanglement error sensitivity

As discussed in Sec 2.4, qubit-qubit entanglement is driven by transiently entangling the internal state with the motional state. Consider a single ion driven by a state-dependent force with spin phase  $\phi_s = 0$ . As shown in Eq. (2.31), the residual qubit-motion entanglement can be described by a unitary displacement  $D(\alpha\sigma_x)$ , where  $\alpha$  quantifies the strength of the residual entanglement. The qubit error caused by  $D(\alpha\sigma_x)$  can be quantified by taking the partial trace over the motional states. For an initial qubit state  $|0\rangle$  we find a bit-flip error

probability of:

$$P(|1\rangle) = 1 - \frac{1}{2} \left( 1 + e^{-2|\alpha|^2} e^{-2|\alpha|^2 \bar{n}} \right).$$


For  $|\alpha| \ll 1$ , we can approximate this error as

$$P(|1\rangle) \approx |\alpha|^2 (1 + 2\bar{n}).$$

In other words, thermal occupation of the motional mode makes it more challenging to entangle and then disentangle with the oscillator. While the process can in principle be perfect, any errors get amplified by a factor of  $(1 + 2\bar{n})$  when performed on a thermal state.

### Mode frequency couplings

In the first approximation, the frequency of every mode is set by the external potential. When multiple ions are trapped in a single harmonic well, their normal modes depend on the competition between the external fields and the ion-ion repulsion. Higher Fock states of radial modes increase the distance between the ions and hence lower their mutual repulsion. This is illustrated in Fig. 2.15. When the radial modes are thermally occupied, this results in a shot-to-shot fluctuation of some normal modes frequencies [Roo08].



**Figure 2.15:** a) Equilibrium spacing of two ions (black line) determines the axial stretch mode frequency  $\omega_s$  b) Higher fock states of the radial breathing mode (red) result in larger mean spacing between the ions (black line) c) As a result, the axial stretch mode oscillator (blue) is coupled to the radial stretch mode oscillator (red).

As an example, consider the axial modes of two trapped ions. When the radial modes are in the ground state, the axial stretch mode frequency is given by  $\omega_s = \sqrt{3} \times \omega_z$ , where  $\omega_z$  is the centre-of-mass frequency of one ion. When a radial mode at a frequency  $\omega_r$  is in Fock state  $n_r$ , the axial stretch mode frequency experiences a *cross-Kerr* frequency shift of  $\delta\omega_s = \chi n_r$ , with<sup>17</sup>:

$$\chi = -\omega_s \left( \frac{1}{2} + \frac{\frac{1}{2}\omega_s^2}{4\omega_r^2 - \omega_s^2} \right) \left( \frac{\omega_z}{\omega_r} \right) \left( \frac{2\hbar\omega_z}{\alpha^2 m c^2} \right)^{1/3} \quad (2.39)$$

where  $\alpha$  is the fine-structure constant and  $m$  is the mass of one ion [Nie08]. When the

<sup>17</sup> Beware: the original paper [Roo08] has a mistake in the formula

## 2. QUANTUM INFORMATION WITH INDIVIDUAL TRAPPED IONS

radial mode is thermally occupied with mean Fock state  $\bar{n}_r$ , this results in a shot-to-shot axial stretch frequency fluctuation of

$$\langle \delta\omega_s^2 \rangle = \chi^2 \bar{n}_r (2\bar{n}_r + 1). \quad (2.40)$$

On the other hand, the axial centre-of-mass frequency is set only by the external potential and is therefore left unaffected by this coupling.

As this example illustrates, certain motional modes exhibit frequency fluctuations when other modes are thermally occupied. This can limit the fidelity of sideband operations (and hence two-qubit gates). However, assuming that the heating and dephasing rates are low enough, the motional frequencies are nonetheless constant over the duration of each experiment, allowing one to suppress this error channel through gate phase-modulation (Sec. 2.4.3).

### 2.5.2. Doppler cooling

Doppler cooling is the first stage of cooling the oscillator towards the ground state. It is implemented by applying a red-detuned laser on a cycling transition. In the classical picture, the Doppler shift increases the scattering cross-section for particles counter-propagating with the laser light, and so emitted photons have larger average energy than the absorbed ones. This leads to an overall cooling of the atom [Ste86].

We can quantitatively analyse the limits of the process by combining the results from Sec. 2.2.5 and Sec. 2.3.4. The Lamb-Dicke approximation does not apply at the start of the cooling, but close to the steady-state when we can think of the process in terms of carrier and first-sideband interactions. Typically the transition linewidth  $\Gamma$  exceeds the mode frequency  $\omega_m$ , and so all three terms in Eq. (2.15) are driven simultaneously. For a lineshape function  $p_0(\delta)$ , this gives:

$$\begin{array}{ll} P_B \propto \eta^2 (n+1) p_0(\delta - \omega_m) & \text{BSB excitation probability, } |n\rangle \rightarrow |n+1\rangle \\ P_C \propto p_0(\delta) & \text{Carrier excitation probability, } |n\rangle \rightarrow |n\rangle \\ P_R \propto \eta^2 n p_0(\delta + \omega_m) & \text{RSB excitation probability, } |n\rangle \rightarrow |n-1\rangle \end{array}$$

By picking  $\delta < 0$  we make the RSB excitation (which takes away a quantum of motion) more likely than the BSB excitation (which adds a quantum of motion). The following

atomic decay branches at a ratio of:

$$\begin{array}{ll}
 P'_B \propto \eta_s^2(n+1) & \text{BSB decay probability, } |n\rangle \rightarrow |n+1\rangle \\
 P'_C \propto 1 & \text{Carrier decay probability, } |n\rangle \rightarrow |n\rangle \\
 P'_R \propto \eta_s^2 n & \text{RSB decay probability, } |n\rangle \rightarrow |n-1\rangle
 \end{array}$$

where the Lamb-Dicke parameter for spontaneous emission  $\eta_s$  is different from  $\eta$  since the outgoing photons do not have a well-defined direction ( $\eta_s^2 = k^2 z_0^2/3$  for isotropic spontaneous emission).

The combined effect of excitation followed by emission at a rate of  $\Gamma$  is that the motional state  $n$  heats at a rate  $R_+ \propto \Gamma(P_C P'_B + P_B P'_C)$  and cools at a rate of  $R_- \propto \Gamma(P_C P'_R + P_R P'_C)$ . The final temperature can be found by equating the heating rate with the cooling rate. With this analysis, we find that the final temperature is the lowest when the effective transition linewidth is the narrowest. For any given transition, this is achieved in the limit of no power broadening, and so  $s_0 \rightarrow 0$ . The lowest temperature is then obtained when  $\delta = \Gamma/2$  and for a two-level system is given by [Esc03; Ste86]:

$$\bar{n} = \frac{\Gamma}{\omega_m} \frac{\cos^2 \theta + \frac{1}{3}}{4 \cos^2 \theta} \quad (2.41)$$

where  $\theta$  is the angle between the cooling laser and the oscillator mode. For  $\theta = 0$ , we find  $\bar{n} = \Gamma/(3\omega_m)$ , while  $\theta = 45^\circ$  gives  $\approx 0.4\Gamma/\omega_m$ . At the angle of  $\theta \approx 65^\circ$ , the Doppler temperature increases by a factor of two from the optimum.


Doppler cooling is typically the longest step of our experiments, although the cooling rate can be increased at the cost of increasing the final temperature. Note that the cooling occurs for all the normal modes in parallel, with the final temperature inversely proportional to the mode frequency. Finally, Eq. (2.41) implies that  $\bar{n}$  diverges for  $\theta \rightarrow \pi/2$ , further highlighting the need to tilt the radial modes in SETs, where the cooling beam typically travels parallel to the chip surface (Sec. 2.1.3). For the  $\lambda = 397$  nm cycling transition and a typical axial frequency  $\omega_m = 2\pi \times 2$  MHz we find  $\Gamma/\omega_m \approx 11$ , limiting the temperature to  $\bar{n} \approx 3.5$  in the two-level approximation. Further cooling is therefore necessary to reach the true oscillator ground state.

### 2.5.3. EIT cooling

As evidenced by Eq. (2.41), the lower the transition linewidth, the lower the final temperature. electromagnetically-induced transparency (EIT) cooling allows us to cool below the Doppler limit by narrowing the excitation linewidth, thus limiting the carrier excitations. In this section, I describe the intuition behind this technique. An interested reader is encouraged to consult [Roo00; Mor00] for detailed mathematical description of the cooling process and its limits.

## 2. QUANTUM INFORMATION WITH INDIVIDUAL TRAPPED IONS

EIT cooling is related to a phenomenon called coherent population trapping [Ari96]. Consider a pair of drives as shown in Fig. 2.16.



**Figure 2.16:**  $\Lambda$  configuration necessary for coherent population trapping. Whenever  $\delta_1 = \delta_2$ , the population gets pumped into a dark state of the qubit, whereby it stops emitting photons.

Whenever  $\delta_1 = \delta_2$ , there is a so-called dark state

$$|\Psi_D\rangle = \frac{\Omega_2}{\Omega}|0\rangle + \frac{\Omega_1}{\Omega}|1\rangle,$$

with  $\Omega = \sqrt{\Omega_1^2 + \Omega_2^2}$  for which the drives destructively interfere, removing any chance of excitation to  $|e\rangle$ . As a result, if  $|e\rangle$  can decay to both ground states, the population will be pumped to  $|\Psi_D\rangle$ , at which point it stops scattering photons. Therefore, when we scan the drive detuning  $\delta_1$ , we will find an excitation suppression (or a “dark resonance”) at  $\delta_1 = \delta_2$ .

This phenomenon is exploited in EIT cooling as follows. A strong  $\sigma$ -polarized “pump” drive is far-detuned by  $\delta_2$  from the atomic transition  $|0\rangle \leftrightarrow |e\rangle$ . This creates a Stark-shifted “bright” dressed state, which can be resonantly excited from  $|1\rangle$  at the frequency of  $\delta_1 = \delta_2 + \Omega_1^2/(4\delta_2)$  by a weak  $\pi$ -polarized “probe” beam (Fig. 2.17). Setting

$$\delta_1 = \delta_2, \quad \text{and} \quad \Omega_2 = 2\sqrt{\omega_m \delta_2} \quad (2.42)$$

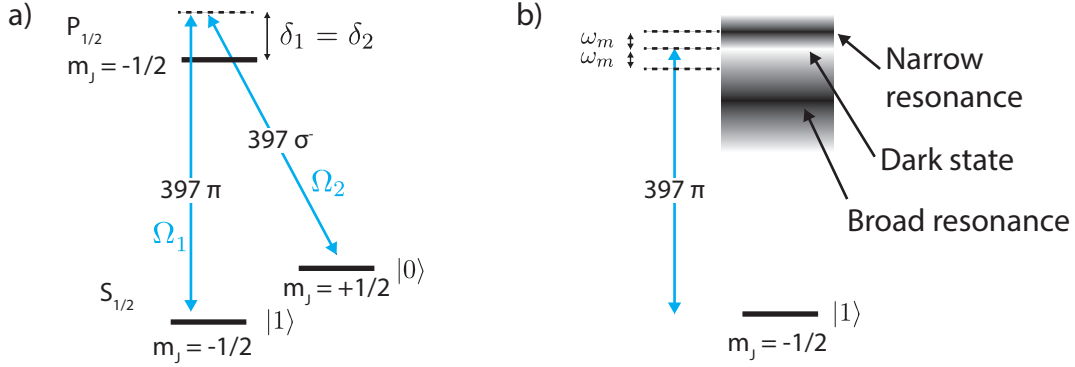
aligns this resonance with the RSB of the “probe” drive. On the other hand, the carrier excitation is completely suppressed, since  $\delta_1 = \delta_2$  is a dark resonance<sup>18</sup> The BSB excitation can be made weak by setting  $\delta_2 \gg \Gamma$ , since the next (wide) resonance is at  $\delta_1 = -\Omega_1^2/(4\delta_2)$ , as shown in Fig. 2.17 b).

Since the carrier excitation is completely eliminated, the final temperature is now predominantly limited by the residual BSB excitation. This results in the cooling limit of [Mor00]:

$$\bar{n} = \left( \frac{\Gamma}{4\delta_2} \right)^2.$$

<sup>18</sup> It turns out that, in the EIT scenario, this dark resonance persists even when the pump and probe fields are not phase coherent [Kha15]. This is in contrast to the standard coherent population trapping.





**Figure 2.17:** EIT cooling configuration. a) Necessary laser transitions. EIT is achieved by combining a strong  $\sigma$ -polarized pump beam with a weak  $\pi$ -polarized probe beam. Both beams must be detuned by the same amount from the excited state to create conditions for coherent population trapping. b) The resonant excitation structure as seen by the probe beam. A dark state at  $\delta_1 = \delta_2$  results from coherent population trapping. Optimal cooling is achieved by aligning the drive RSB with the narrow resonance.


Thanks to this relationship, the cooling is only limited by how much the lasers can be detuned, which is a technical limitation and not a fundamental one. We usually operate with  $\delta \approx 2\pi \times 50$  MHz, which should be sufficient to cool a  $^{40}\text{Ca}^+$  ion to  $\bar{n} \approx 0.01$ . We do not reach such low temperatures in practice, likely due to polarization impurity and quantization axis oscillations (Sec. 5.2.1). Nonetheless, EIT does allow us to cool close to the motional ground state (Sec. 5.6.3). EIT can cool multiple modes in parallel since the carrier suppression is power-independent, but the optimal AC Stark shift in Eq. (2.42) can only be achieved for one mode [Lec16]. The cooling rates are proportional to the overlap between the oscillator mode and the wavevector difference of the “pump” and “probe” beams.

#### 2.5.4. Resolved sideband cooling

A different way of suppressing the carrier and BSB excitations is to use resolved sideband cooling. Each cycle of sideband cooling consists of two steps, labelled (A) and (B) in Fig. 2.18. Consider an ion prepared in the internal state  $|1\rangle$  and Fock state  $|n\rangle$ . In the step (A), the ion is excited on the RSB of a narrow-linewidth optical qubit transition, changing the combined ion-oscillator state to  $|0\rangle \otimes |n-1\rangle$ . Afterwards in step (B), the ion’s internal state is reset to  $|1\rangle$  by optical pumping on a dipole transition (Sec. 2.9). Thus, at the end of the cycle, the ion’s internal state is returned to  $|1\rangle$ , but the motional state is reduced to  $|n-1\rangle$ . As a consequence, after sufficiently many cycles, the ion is pumped into the motional ground state, regardless of the input state. Since different Fock states  $n$  correspond to different RSB Rabi frequencies  $\Omega\eta\sqrt{n}$ , one cannot tailor the RSB drive to be an exact pi-pulse for a thermal state of motion. To make the cooling most efficient, we

## 2. QUANTUM INFORMATION WITH INDIVIDUAL TRAPPED IONS

ramp the pulse time from an initial value of  $t \approx \pi/\Omega\eta\sqrt{\bar{n}}$  to the final value of  $t \approx \pi/\Omega\eta$ .



**Figure 2.18:** Resolved sideband cooling is composed of two steps. Drive (A) is an RSB drive on an optical qubit transition (any transition can be used for this surface). The decay step (B) is stimulated by applying the repump lasers. After applying (A) and (B) in sequence, the motional occupation reduces until it reaches  $n = 0$ .

In the leading order, the limit to the temperatures achievable with resolved sideband cooling comes from off-resonant carrier and BSB excitations (Sec. 2.3.3). For square pulses, the final temperature is given by [Ita]:

$$\bar{n} \approx \frac{\Omega^2}{\omega_m^2} \left( \eta_s^2 + \frac{1}{4} \eta^2 \right).$$

Since off-resonant excitations be reduced by pulse shaping or by decreasing  $\Omega$ , the practical limit on  $\bar{n}$  is set by how long we're willing to spend sideband cooling. However, the Rabi frequencies for quadrupole transitions are lower than for dipole transitions, and thus sideband cooling is naturally slower than Doppler or EIT cooling. When slowed down further, the achievable temperatures might become limited by mode heating rates.

Unlike the other cooling methods, sideband cooling is single-mode. When the quadrupole beam is global, only one mode can be cooled per round, and cooling  $M$  motional modes requires  $M$  separate sub-cycles, each at a different RSB frequency<sup>19</sup> [Mor99]. Therefore, the sideband cooling time grows linearly with the number of ions  $N$ . However, if  $N$  ions in a chain are individually addressed, it is possible to perform  $N$  cooling cycles in parallel, and thus keep the cooling time approximately constant<sup>20</sup> [Che20].

<sup>19</sup> In our experiment, we sideband-cool the motional modes with the highest heating rate last to achieve the lowest average temperatures

<sup>20</sup> Assuming the amount of available laser power grows  $\propto N$

### 3. Experimental setup

Trapped-ion quantum information experiments require complex setups combining laser optics, electronics and vacuum systems. Our setup was designed and built to operate planar Paul SETs in a cryogenic environment. The trap is housed inside a copper chamber, which is cooled down to 4 K using a closed-cycle cryocooler. The chamber is placed under vacuum, and heat-shields protect it from excess room-temperature black body radiation. Stable trapping is achieved by a combination of DC and RF voltages applied to trap electrodes. The 4 K chamber houses an RF amplifier, “low-pass” RC filters, as well as various ancillary electronics. The light emitted by the ion is collected by a cryogenic objective and focused onto a photo-multiplier tube (PMT) and/or a camera outside the vacuum chamber. Quantum control is achieved primarily through laser pulses, delivered either free-space or via integrated optics.

The experimental system underwent significant changes during my time in the group. This was motivated by the obstacles and limitations we encountered in the original setup built by F.M. Leupold and J. Alonso, henceforth dubbed *setup 1.0*, which is described extensively in [Leu15]. As a result, different experiments were performed in different setups and traps, as shown in the timeline Fig. 3.1. In this chapter, I present the updated cryogenic setup 2.0, built together with C. Zhang, with significant help from T.-L. Nguyen and J. Alonso.

#### Original setup

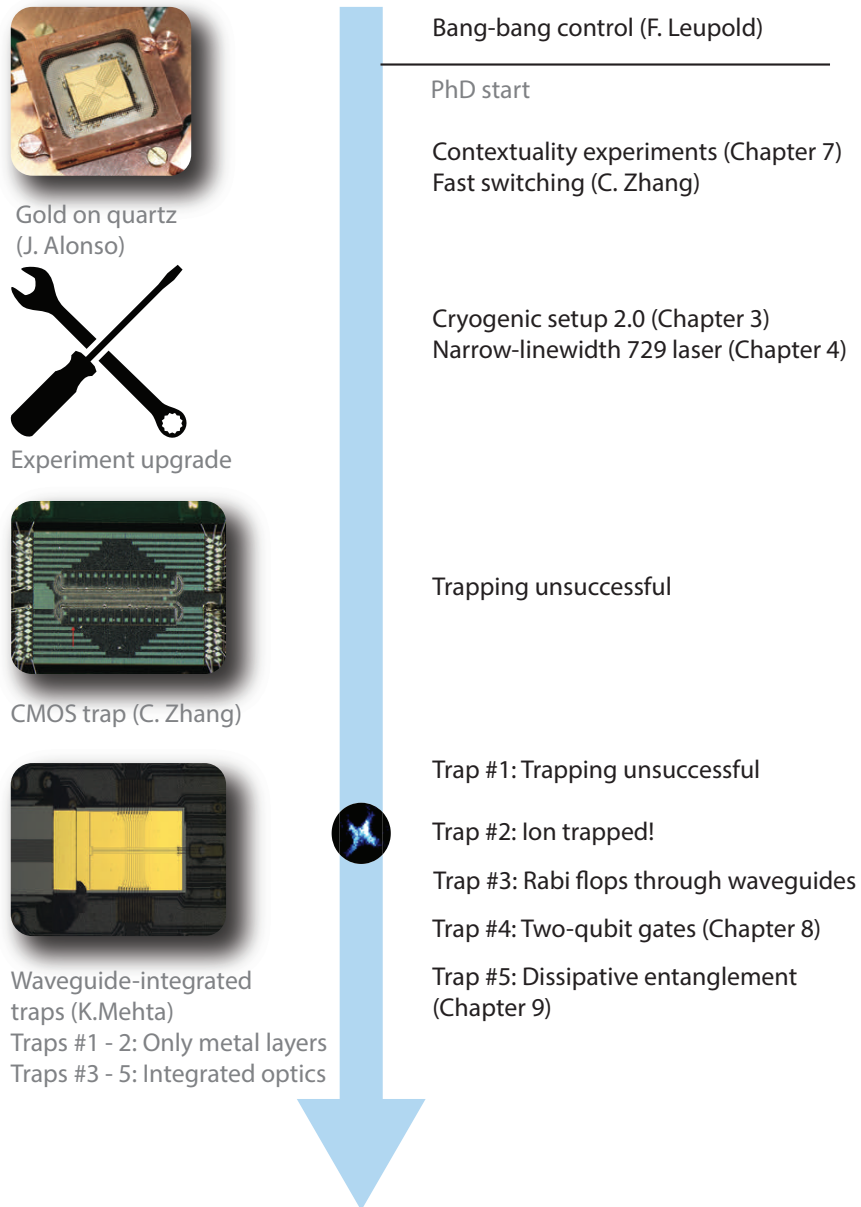
In every section, I discuss the strengths and weaknesses of the original design, and how they influenced the upgrade process. To clearly separate the old elements from current ones, the original setup is presented in boxes like this one.

Sec. 3.1 discusses why we use cryogenic environments, and how the cryogenic vacuum system was reconstructed. Sec. 3.2 focuses on the 4K chamber design. In Sec. 3.3 we describe all the ion traps tested in this thesis. Finally, Sec. 3.4 describes the improvements to the electrical control and the magnetic field stabilisation. Since I spent a comparable amount of time on the laser setup as I did on all the other parts of the experimental setup combined, the optics and light delivery is discussed separately in Chapter 4.

#### 3.1. Cryogenic vacuum system

Traditionally, ion-trap experiments were conducted in room-temperature (RT) UHV systems. In Sec. 2.1.3 we highlighted the difficulties created by miniaturised SETs. A cryogenic

### 3. EXPERIMENTAL SETUP



**Figure 3.1:** Timeline of my PhD.

vacuum environment allows us to fight many of these problems by:

1. **Reducing heating rates.** Anomalous heating can be suppressed by using a cryogenic environment [Lab07; Des06]. This is thought to be because a lot of the conjectured mechanisms involve thermally activated noise sources [Bro14].
2. **Improving ion lifetime.** Ion lifetime is determined by the trap depth and the frequency and energy of background gas collisions. Cryogenic temperatures allow for the generation of ultra-low pressures, and the remaining background gas molecules have a lower energy [Pag19; Ant09]. This extends the lifetime of a single laser-cooled ion in a typical SET from a few hours to essentially forever<sup>1</sup>.
3. **Increasing reconfigurability.** Unlike RT chambers, cryogenic systems do not require baking to achieve UHV. As a result, they allow for easy changes of in-vacuum instruments, such as trap chips or electronics. This makes them perfect for testing novel components [Stu18]. Furthermore, material constraints are also relaxed compared to RT experiments.

However, low temperatures create a number of challenges and constraints:

1. **Vibrations.** Cryogenic cooling generally causes additional vibrations of the trap. While the strength and spectrum of vibrations strongly depend on the type of cryostat and how it's mounted, typical displacements are  $\approx 1 \mu\text{m}$  for wet cryostats and  $\approx 10 \mu\text{m}$  for closed-cycle units [Bra16a]. Vibrations in a closed-cycle system can be reduced to nanometre-scale only through careful structure design and vibration isolation [Mic19].
2. **Limiting heat load.** Maintaining low temperatures requires limiting the thermal load on the 4K stage. While black-body radiation is effectively suppressed by heat shields, care must be taken to minimise heat flow through the cables, as well as thermal dissipation from passive and active electrical components.
3. **Temperature compatibility of components.** A lot of standard electrical, optical and mechanical components do not function properly in cryogenic environments. For example, thermal contraction distorts tightly toleranced components (such as imaging objectives) and doped semiconductors change their behaviour at low temperatures<sup>2</sup>, causing some active electronics to malfunction.

Our cryostat is a closed-cycle pulse tube (PT) cryocooler<sup>3</sup>. The 2nd stage reaches the temperature of 4 K with 0.9 W cooling power, while the 1st stage is specified for 31.5 W at 45 K. PT cryocoolers generally suffer from considerable vibrations induced by the flow of helium in and out of the PT, at a frequency of 1.4 Hz. In 2010, engineers at Cryomech

---


<sup>1</sup> In our experiments, single-ion lifetime is limited by how often lasers come out of lock

<sup>2</sup> Due to so-called *carrier freeze-out*

<sup>3</sup> PT410-RM from Cryomech

### 3. EXPERIMENTAL SETUP

showed that they can reduce 2nd stage sample vibrations to below 0.5 nm by opening a gap between the PT and the sample and using helium as heat exchanger gas [Wan10]. The associated cooling power decreased slightly, to 0.7 W at the 2nd stage and 30 W at the 1st stage. While reducing cryogenic vibrations is important, the figure of merit for ion trapping is not absolute vibrations, but rather differential vibrations between the ion and the laser beams. In [Wan10], the PT was rigidly attached to the table, while the cryogenic sample was freely suspended from the bellows, as shown in Fig. 3.2 (a).



**Figure 3.2:** a) Schematic of the cryostat in [Wan10]. The PT is rigidly mounted onto the table. The vacuum chamber, heat-shield and cryogenic chamber are suspended from steel bellows. b) Schematic of the cryostat initially used in our lab. The vacuum chamber is rigidly mounted onto the table, hence fixing the position of the cryogenic chamber. The PT is floating on steel bellows.

If used in our experiment, it would have meant that the position of the ion trap was not directly referenced to the optical table, raising a number of possible issues regarding laser alignment. Therefore, in the setup 1.0, a decision was made to attach the vacuum chamber rigidly to the optical table, as shown in Fig. 3.2 (b). It meant that the PT was directly floating on the steel bellows. While this setup allowed successful trapping, it had several disadvantages.

#### Cryogenic vacuum system 1.0

##### Weakness: Vibration issues

It rapidly became clear that the PT, when freely floating on bellows, induced

significant ( $\approx 10\ \mu\text{m}$ ) vibrations of the ion trap relative to the optical table. This happened for three reasons:

1. The PT is relatively light. Therefore, even small momentum kicks from the helium flow cause significant displacements and, in turn, the 1st stage heat exchanger collides with the shield cooling station at every compression cycle.
2. The PT connects to multiple helium lines. These are positioned asymmetrically, tilting the PT and creating a permanent contact between the 1st stage heat exchanger and the shield cooling station.
3. The surface on an air-floated optical table moves significantly in response to momentum kicks.

The initial solution was to create an additional rigid attachment between the PT and the vacuum chamber using three posts with adjustable heights. We would then change the tilt while roughly measuring vibrations by monitoring the transmission of a laser beam focused onto the trap. This was a very long and tedious process. Furthermore, while it was possible to significantly decrease vibrations, their impact on coherent operations was clearly visible and was one of the dominant decoherence mechanisms. The reason was two-fold. First of all, this attachment eliminated the original premise of the mechanical decoupling, allowing PT vibrations to propagate to both the ion trap and the optical table. Secondly, as I shall discuss in Sec. 3.1.1, it is likely we never managed to align the PT well enough to fully avoid mechanical contact.

#### Weakness: Assembly issues

While in principle reconfigurable, the ability to open and modify the cryogenic system was hindered by the optical table, which blocked the path to sliding the vacuum chamber and the heat shield on and off. In order to access the insides of the setup, the whole cryostat, including its mounting structure and the helium lines, had to be lifted off the optical table with a crane. In addition to making opening up time-consuming and laborious, this prevented any build-up of permanent optical or mechanical structures around the experiment (e.g. frames to isolate vibrations by attaching the PT to the ceiling).


#### Weakness: Beam pointing drifts

While mounting the vacuum chamber directly to the optical table helped reduce drifts, it did not eliminate them completely. One dominant issue was that any changes to the 4K chamber temperature resulted in visible trap displacements. Most annoyingly, after running the atomic oven (located inside the 4K chamber) for a few minutes, one had to wait a few more minutes for the beam pointing to return to the original state.

### 3. EXPERIMENTAL SETUP

While we suspected this to be caused by thermal expansion of mechanical parts, it turned out to be caused, almost entirely, by the pressure change of the helium buffer gas.

The remainder of this section describes how the cryogenic vacuum chamber 2.0 was designed and built to eliminate the problems highlighted in the box above. The schematic of the redesigned cryostat is shown in Fig. 3.3. It addresses the vibration issues by rigidly mounting the PT to the ceiling through a 3D translation platform, and by increasing the clearance between the PT and the dewar walls. Opening and closing are significantly easier, thanks to a two-piece construction of the vacuum chamber and the heat shield. The beam pointing drifts are largely eliminated by dewar pressure regulation, and by placing the oven at the room-temperature stage.



**Figure 3.3:** Schematic of the updated cryostat design.

#### 3.1.1. Pulse tube stabilisation

One of the big unknowns of the setup 1.0 was the exact clearance between the PT and the dewar. In early 2018, I took out the PT and measured the gap to be a mere 0.3 mm over a length of 3 cm, insufficient to ensure mechanical decoupling (the alignment is done by tilting the top of the PT approximately 20 cm away). We decided to enlarge this spacing at



the expense of reduced cooling power<sup>4</sup>. The PT diameter was decreased on a lathe from 129.5 mm to 127.0 mm, widening the gap to 1.6 mm. The PT was subsequently tested for performance in vacuum, and the 2nd stage heat exchanger was buffed to remove observed corrosion<sup>5</sup>. All the adjustments were performed by Cryomech, and the modified PT was mounted back in the dewar. Afterwards, Cryomech sent us foam that they forgot to re-attach into the PT. Since taking the PT out is a very disruptive and delicate process, we have not re-installed the foam<sup>6</sup>.

In order to isolate its vibrations, we attached the PT to the ceiling. We constructed an alignment platform<sup>7</sup> shown in Fig. 3.4 which allows us to position and tilt the PT before clamping it in place. The PT is attached to a 10 mm thick stainless steel (SS) top plate with four M6 bolts. An 80 deg cut-out in the top plate provides clearance for the helium inlet tube and electrical connections of the PT (not shown). The tilt is adjusted using three M10 tilt screws at a radius of 170 mm. Afterwards, the top plate is fixed to the middle plate using three locking bolts. The middle and bottom plates are both 6 mm thick SS squares, with a cut-out that allows them to be inserted from the side. They are separated by 4 mm rubber spacers, which allow the middle plate to slide when pushed by M8 bolts housed inside aluminium blocks attached to the bottom plate. The middle plate can be translated by  $\pm 10$  mm in two horizontal directions and is afterwards fastened to the bottom plate using four bolts and (optionally) table clamps. The bottom plate is attached to the ceiling with four 50 mm extruded aluminium profiles and square brackets.

The alignment of the PT was performed with an accelerometer attached to the cold plate<sup>8</sup>. Vibrations were induced either by a remote motor or artificially with a computer fan attached to the PT. We then moved the position of the top and middle plates until the measured vibrations were minimised. At those positions, we do not record any transmission of vibrations from the PT onto the cold plate.

Unfortunately, even in the absence of PT perturbations, the accelerometer records movements of  $> 100$  nm, which prevent us from measuring the vibrations with an interferometer. We observe structural resonances at 99 Hz, 139 Hz and 310 Hz, and find that the low-frequency hum of the air conditioning clearly enhances the vibrations<sup>9</sup>. Despite improvements to the mounting structure (such as in Fig. 3.5), these could not be completely eliminated. Compared to other similar cryogenic systems in our laboratory<sup>10</sup>, the difference

---

<sup>4</sup> We did not receive information from Cryomech about the exact relation between the two, so we had to guess and hope for the best

<sup>5</sup> The speculation is that this was caused by helium impurity. Our heat exchanger helium is labelled as 99.996% pure, but perhaps this does not reflect the reality.


<sup>6</sup> The exact role and impact of the foam remains mysterious to me. The manufacturer's reply when asked about the purpose of the foam was: "The foam can help to improve the cooling a little. The system will work without it"

<sup>7</sup> Based on a design by R. Matt

<sup>8</sup> In systems with rubber bellows, it is possible to check for a collision between the PT and the dewar electrically. However, since our bellows are SS, this method was not successful

<sup>9</sup> There are possibly more acoustic resonances at lower frequencies, but we did not have an appropriate signal generator to excite them

### 3. EXPERIMENTAL SETUP



**Figure 3.4:** PT mounting assembly rendering. Top plate controls the tilt, the middle plate controls the horizontal position, and the bottom plate mounts directly to the ceiling.


is that our helium dewar is simply less rigid. Gently pushing the 4K chamber with a finger results in significant oscillations, which begin at low frequency, and become faster before dissipating a few seconds later. It seems that our cryostat efficiently picks up the environmental noise, leading to pronounced vibrations even with the cryocooler switched off. This might be further influenced by the stiff SS bellows – the other low-vibration cryostats in our laboratory are fitted with rubber ones.

#### 3.1.2. Helium dewar pressure control

In the absence of mechanical contact, heat is transferred to the PT via helium buffer gas. We have discovered that the change of helium pressure is a dominant cause of trap chip displacement when cryogenic thermal load changes. To stabilise the trap position, we have installed a simple helium pressure regulation system. When the temperature rises, the helium pressure inside the dewar increases until  $p_{\max}$ , when a pressure relief valve (PRV)

---

<sup>10</sup> Based on the closed-cycle ColdEdge cryostat with the Ultra-Low Vibration interface



**Figure 3.5:** The last stage of the mounting structure improvement was the clamping jack, which rigidly attaches the bottom of the vacuum chamber to the optical table. While the structure does not directly attach to the experiment, nonetheless it decreases the trap vibration amplitude by  $\approx 20\%$ , presumably by eliminating a resonance of a freely oscillating vacuum chamber.

opens the flow to the building recovery line. When the temperature decreases and the pressure falls below a threshold value  $p_{\min}$ , a regulator opens a valve from a helium bottle to keep the pressure constant.


Before filling with helium, we evacuate the air from the dewar and its circuitry using a turbo-molecular pump (TMP). At the pressure of  $1 \times 10^{-5}$  mbar the pump is disconnected and the regulator is set to  $p_{\min} \approx 1.8$  psig, slightly below the PRV opening pressure  $p_{\max} \approx 2.1$  psig. Figure 3.6 shows dewar pressure over a typical work day. Within this pressure range, we do not observe trap position drifts.

In July 2017, during a cryostat cool-down, a leak opened on the helium dewar. We first noticed that, even at room temperature, the vacuum pressure in the main chamber quickly rose when the dewar was filled with helium and decreased when the dewar was evacuated. We then removed the vacuum chamber and located the leaks directly by keeping the dewar helium pressure at 2 psig and using a helium leak detector in a “sniffer” mode. We identified a total of three leaks on the dewar brazes. The leaks were sealed<sup>11</sup>, but it is an ongoing concern that they may open at some point.

---

<sup>11</sup> using Stycast 2850FT

### 3. EXPERIMENTAL SETUP



**Figure 3.6:** Helium pressure over a typical work day. Sharp pressure rises mark loading events, when the oven increases the radiative heatload. At the pressure of 2.1 psig the outlet valve opens, and once the oven is switched off, the pressure returns to equilibrium

#### 3.1.3. Modular radiation shields

In order to enable access to the 4 K chamber without removing the setup from the optical table, a single copper radiation shield was replaced with two aluminium shields shown in Figure 3.7a. The top shield is attached to the shield cooling station with brass screws and spring washers. To ensure maximum heat transfer, the bottom shield is tightly fastened to the top shield with twelve steel M5 bolts through a layer of Apiezon<sup>12</sup>.


The bottom shield features ten openings: on the surface of the cylinder there are eight 40 mm diameter holes for laser and fibre access, and a 8 mm hole for neutral atoms. A 35 mm diameter hole on the end of the cylinder is used for imaging ions from the bottom. Since these cause significant heat load on the 2nd stage, we cover the openings with aluminium flaps, in which we open smaller holes, tailored to whatever laser configuration we want to use.

Under typical trapping conditions, the shield cooling stage thermalises to  $\approx 57$  K. This is significantly higher than  $\approx 48$  K recorded in setup 1.0, and probably reflects reduced cooling power due to decreased PT diameter. We also recorded a shield temperature increase of  $\approx 3$  K when the PT was aligned for minimum mechanical contact. The bottom part of the lower heat shield is  $\approx 6$  K warmer than the top of the upper heat shield, indicating a good thermal contact between the pieces.

After the new radiation shields were manufactured, they did not fit the cryostat. It turned out that the helium dewar is noticeably<sup>13</sup> crooked, resulting in  $\approx 0.6$  deg tilt of the heat shields with respect to the vacuum chamber. This caused a near contact of the shield joining ring with the room temperature vacuum chamber (Figure 3.7b). We have since enlarged the holes on the top radiation shield and used this added degree of freedom to gain extra clearance.

<sup>12</sup> Cryogenic high vacuum grease Apiezon N from Apiezon

<sup>13</sup> Physicists tend to notice it once you point it out. Mechanical engineers notice it immediately



**Figure 3.7:** (a) Two-piece radiation shield. (b) Bottom view of the cryostat during assembly. The 4K mounting plate (copper, middle) is surrounded by the radiation shield joining ring and by the upper vacuum chamber. The shield joining ring is well centered with the 4K mounting plate, but nearly contacts the vacuum chamber in the bottom-right corner of the image.

#### 3.1.4. Two-piece outer vacuum chamber


A new, two-piece vacuum chamber was made to house the experiment. It is made out of SS 304 pieces connected to each other, and to the top cryostat plate, with rubber seals. On the top plate, we cut off and permanently sealed the four existing flanges to make room for the vibration isolation assembly. Cables are now routed through six KF flanges<sup>14</sup> in the top chamber. The bottom chamber features nine DN40 CF flanges (eight on the sides, one on the bottom) and one DN16 flange. The unusual choice of CF flanges on the bottom was made in the hope of avoiding the 4K viewport coating we periodically observed in the cryostat 1.0. Indeed, we never observed any viewport deterioration in cryostat 2.0, though an increased viewport temperature (Sec. 3.2.3) could also explain this.

#### 3.1.5. Room-temperature oven

In order to facilitate loading without imposing a high heat load on the cryogenic stage, we mounted a resistively heated oven on the vacuum flange. The discussion of oven design and cryogenic atomic sources is presented in Sec. 3.2.5. Here, I just show the key parameters of the oven, which is shown in Fig. 3.8.


<sup>14</sup> DN40KF

### 3. EXPERIMENTAL SETUP



**Figure 3.8:** Flange-mounted oven before assembly.

The resistively heated oven consists of a 1.27 mm inner diameter SS316 tube with tantalum wires spot-welded to either end. It is held in place by copper wires, which are themselves clamped through a Macor spacer on an aluminium L-bracket. The oven itself is initially mounted in a Macor spacer to centre it on the flange. Once the copper wires are clamped, the spacer is removed in order to reduce thermal losses. The oven aperture is  $\approx 135$  mm away from the trap, with apertures in between used to collimate the neutral plume (Fig. 3.9).



**Figure 3.9:** Schematic of the room-temperature oven layout. The first aperture is on the heat shield, and it shields the 4 K stage from excess black-body radiation of the oven.

Due to the aforementioned manufacturing misalignments of the cryostat, the centre of the vacuum flange is offset from the optical axis by a few millimetres. To align the oven to the aperture of the trap chip, we mount the flange on a port aligner<sup>15</sup>. The port aligner is first roughly adjusted by looking through the opposing viewport. Once under vacuum, its position is further adjusted by maximising the neutral fluorescence observed with a  $\lambda = 423$  nm beam aligned at the trap centre. During the initial alignment, we use a current of 5.5 A, which creates a very large neutral fluorescence signal of  $\approx 5000$  counts in a  $\tau = 250$   $\mu$ s detection window, visible also to the naked eye. For loading, we operate the oven at 3.6 A, where the neutral fluorescence is just about visible on the PMT. The

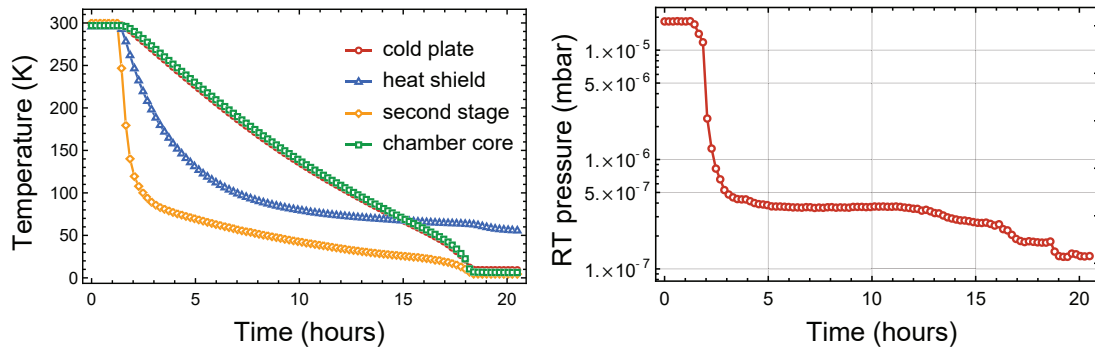
<sup>15</sup> PA35-H from Lesker



alignment procedure needs to be repeated after opening and closing the system, which is annoying but not very difficult.

### 3.1.6. Pumping and cool-down

The closed chamber is pumped with a roughing pump and a TMP until the pressure of  $\approx 2 \times 10^{-5}$  mbar. At that point, we switch on the cryostat compressor and begin the system cool-down. A typical trace of recorded temperatures and pressures is shown in Fig. 3.10. The system reaches the final temperature within approximately 20 hours. As



**Figure 3.10:** Logs of a typical cooldown. (Left) Temperatures of different elements. The heat shield sensor is mounted on the 1st stage heat exchanger. (Right) Pressured measured at a room-temperature gauge.

the pressure drops below  $1 \times 10^{-6}$  mbar, the TMP is no longer useful and is disconnected to avoid vacuum degradation. In the setup 1.0, the chamber would cool down to 4.5 K with the RF on, while the base temperature measured near the trap in the current setup is 6.0 K, likely due to the aforementioned PT and heat-shield changes. We continue to refer to the cryogenic chamber as the “4K chamber”, even though most of the experiments were performed at  $\approx 6$  K.

Note that the base pressure reading of  $\approx 1 \times 10^{-7}$  mbar does not represent the pressure in the UHV chamber, but only the residual effect of the cryo-pumping at room temperature. Furthermore, starting approximately a year ago (late 2019), we began recording periodic pressure spikes of  $\approx 3 \times 10^{-6}$  mbar every one hour. We have not found any correlation between the spikes and the ion lifetime.

## 3.2. 4 K chamber

At the centre of the cryogenic vacuum system is the 4 K chamber. It houses the ion trap, electronics (electrode filters and RF resonator) and the imaging objective. The thick and cold copper walls of the chamber attenuate magnetic field fluctuations (Sec. 3.2.1) and provide UHV environment for the ions. The experiments described in Chapter 7 were performed

### 3. EXPERIMENTAL SETUP

using a trap housed in chamber 1.0. Following that work, we decided to start using more complicated traps. This necessitated a chamber re-design to improve its reconfigurability, practicality and performance.

#### 4K chamber 1.0

The original setup has a number of strong points we aimed to replicate, as well as several drawbacks to improve upon:

##### Strength: Magnetic field shielding.

At cryogenic temperatures, copper is extremely conductive. The middle part of the chamber was a single, hollowed-out high-purity copper block. This meant that electrical current could run in a loop of very low resistance, without encountering mechanical joints. Those 2 cm-thick walls effectively cancelled oscillating external magnetic field noise (Sec. 3.2.1), resulting in long ( $\approx 10$  ms) coherence times of Zeeman qubits.

##### Weakness: Lack of modularity

The chamber featured only three feedthrough slots. When we decided to include additional connections (for a heater wire and temperature sensors), we would simply pull the cables out through a large open pinch-off flange, which raised vacuum concerns. Furthermore, a large helical resonator was taking up a majority of the usable space inside the chamber, making it challenging to add or reconfigure components.

##### Weakness: Cabling and assembly

This original process of chamber assembly was tedious and prone to errors. It was difficult to ensure that cables didn't block laser access or short together. The "card edge" connectors (HCC08DREN, HCC20DREN from Sullins) delivering the DC voltages were very difficult to slide onto the filterboard, resulting in a tedious assembly. The force applied to those cables put a strain on the feedthroughs, which required re-soldering numerous times. Finally, soldering cables meant the chamber could not be kept as clean as it ideally would be.

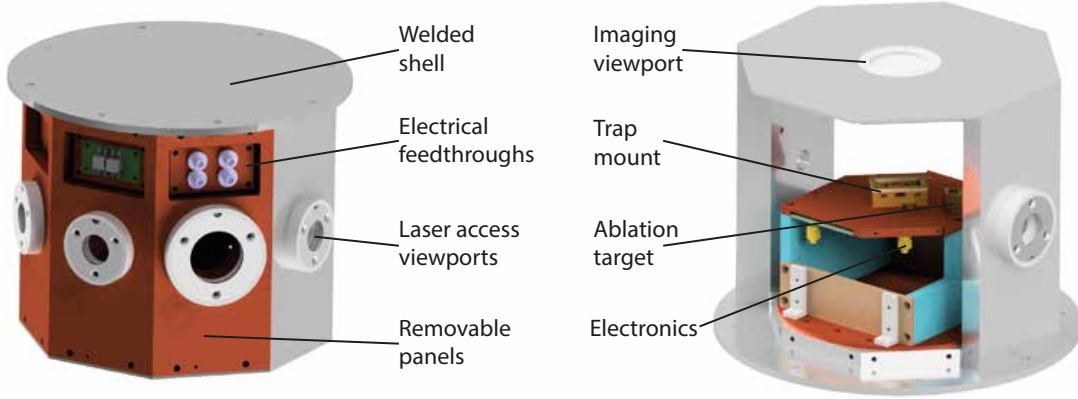
##### Weakness: Oven heat dissipation

The only atom source in the original setup was an effusive oven placed next to the trap. As a result, every loading event produced a significant thermal load on the cryogenic chamber, increasing its temperature and spoiling the vacuum.

The re-designed chamber, shown in Fig. 3.11, improves upon the issues highlighted in the box. The assembly and reconfigurability are made dramatically easier with an open and modular layout. Removable side panels provide easy access, tidier wiring and signal filtering at the same time. Heatable viewports eliminate the need for periodic warm-ups. In addition, changes to electrical components (discussed in Sec. 3.4.1) allow us to place more



elements and improve resonator impedance matching.



**Figure 3.11:** 4 K chamber 2.0. (left) Chamber with closed panels. Removable panels are coloured brown for clarity - in reality, the whole chamber is silver-coated. (right) Chamber placed upside down with panels removed. Copper mouting structure is shown in false colour for clarity. The imaging objective and its piezo stack are omitted for clarity.

### 3.2.1. Magnetic field shielding

A key strength of cryogenic ion trap systems is the reduction in magnetic field noise [Bro11]. A conductive metal chamber will shield its contents from an oscillating B-field by allowing the flow of induced eddy currents [Cel08]. While the exact calculation of the expected attenuation is complicated and geometry-specific, we can use analytical formulae for a simple geometry to guide our intuition. Consider a spherical shell of radius  $r_0$  and thickness  $\Delta \ll r_0$  made out of a material with conductivity  $\sigma$  and relative magnetic permeability  $\mu_r$ . The shell is placed in a uniform external magnetic field  $B_{\text{ext}}$  oscillating at a frequency  $\omega$ . The *skin depth* of the material is defined as

$$\delta = \sqrt{\frac{2}{\omega \mu_0 \mu_r \sigma}}$$

where  $\mu_0 = 4\pi \times 10^{-7} \text{ Hm}^{-1}$  is the vacuum permeability. In the limit of  $\Delta \ll \delta$ , an exact calculation gives the field inside the chamber [Cel08]:

$$B_{\text{int}} = B_{\text{ext}} \frac{3\sqrt{2}\mu_r\delta}{r_0} e^{-\Delta/\delta}.$$

The key takeaway is that the external field is strongly attenuated when the shell thickness is much larger than the skin depth ( $\Delta \ll \delta$ ). Furthermore, attenuation can be improved by using a large shield or low relative permeability.


Many of the main sources of magnetic field noise in trapped ion experiments are synchronised to the socket power line. Therefore, it is beneficial for the magnetic field attenuation to be

### 3. EXPERIMENTAL SETUP

effective already at 50 Hz. At room temperature, the conductivity of oxygen-free electronic (OFE) copper  $\sigma = 5.86 \times 10^7 \Omega^{-1}\text{m}^{-1}$  corresponds to a skin depth of  $\delta \approx 9.3 \text{ mm}$  at  $\omega = 2\pi \times 50 \text{ Hz}$ . It is not realistic to shield the experiment with walls orders of magnitude thicker than  $\delta$ . Therefore, we cannot use such a simple shield to cancel the mains noise effectively. However, the situation changes considerably at cryogenic temperatures, where metal conductivity increases by a factor called the residual-resistance ratio (RRR), which can be as high as 1000 for high-purity copper. For copper with RRR of 150 (realistic for off-the-shelf OFE copper plates [Blo14]), we find the cryogenic skin depth of  $\delta = 0.75 \text{ mm}$ . Therefore, by enclosing the experiment in a chamber with  $\Delta = 20 \text{ mm}$ , we can expect to significantly attenuate magnetic field fluctuations at frequencies of 50 Hz and higher.

#### 3.2.2. Welded shell


The calculations presented above assumed a monolithic shield with bulk conductivity. However, a fully enclosed design allows no access. Once we allow the chamber to be disassembled, its effective conductivity decreases due to finite contact resistance between the bolted parts. This can impair the ability of the system to shield magnetic fields. Our solution, which aims to strike a trade-off between shielding and modularity, is shown in Fig. 3.12.



**Figure 3.12:** 4 K chamber 2.0. (left) Welded cryogenic chamber shell. The top plate attaches to the cold plate, and the bottom viewport is used for imaging. (right) Chamber with side panels can be inserted into the remaining six sides. While the removable panels are shown in a different colour, in reality, the whole chamber is silver-coated gold.

A monolithic shell is formed by the bottom, top and two side panels. This loop is oriented perpendicular to the quantisation axis for the most effective B-field shielding. The remaining six panels are removable, which allows for easy access to the contents of the chamber. While the shell can be created by milling out a single high-purity copper block [Bra16a; Leu15], such a solution is costly and time-consuming. An alternative method is to weld individual panels to form the shell. However, this fusion must be performed under vacuum, to avoid polluting the joint with oxygen, and over a large depth, to ensure

maximum conductivity. An ideal method for achieving this is electron-beam (EB) welding. In this process, electrons are focused on the interface between the two pieces. Upon impact, the material melts, fusing the pieces. The whole process is performed in vacuum, ensuring a clean and oxygen-free interface. Fig 3.13 shows the welded shell prior to silver-coating.



**Figure 3.13:** 4 K chamber 2.0 after welding, prior to silver-coating and cleaning.

The top and bottom lid are 14 mm thick. Inside, numerous screw holes allow for attaching side panels and flexible mounting of chamber contents to the top lid. The bottom plate has holes for mounting an Attocube piezo stack, a charcoal getter and an imaging viewport with 37.1 mm diameter. Both plates feature many more screw holes than currently necessary, such that the experiment can be flexibly modified in the future. The two side panels are 20 mm thick and each contains a single  $\phi = 12.7$  mm hole for laser access. All the pieces are manufactured in-house out of OFE copper<sup>16</sup> without a specified RRR<sup>17</sup>. The four pieces are then EB-welded together<sup>18</sup>. By welding and then cutting a test piece, we verified fusion occurred over the whole 20 mm of the contact area. Following that, the chamber was silver-coated<sup>19</sup> to reduce emissivity and improve thermal contact. A  $\approx 5\mu\text{m}$ -thick layer of silver was applied onto the copper directly, without the usual nickel layer, to avoid unwanted magnetisability. A final inspection revealed that the top and bottom lid were slightly deformed during the EB welding process. We then had to sand some parts to allow side panels to slide in.

<sup>16</sup> C10100, > 99.99% purity copper from Px Precimet SA

<sup>17</sup> In hindsight, that was a mistake

<sup>18</sup> Swissbeam

<sup>19</sup> Work done externally by Doerr AG in Zurich

### 3. EXPERIMENTAL SETUP


#### 3.2.3. Chamber panels

Six out of eight panels of the 4 K chamber are removable. Each silver-coated OFE copper panel features a cut-out for mounting feedthroughs and a 23 mm diameter hole for laser access<sup>20</sup>. Each feedthrough is mounted in a copper plate of an identical size<sup>21</sup>, such that connections can be swapped and modified without replacing the chamber panels. Feedthroughs are recessed into the panel to avoid the cables protruding. The feedthroughs we currently use are described in Sec. 3.4.1. Unused feedthrough cut-outs are closed with blank copper flanges.

##### Unfiltered cables

In cryostat 1.0, only the DC and RF lines were filtered close to the trap. Otherwise, there were a number of unfiltered connections - for the oven, trap heater and the attocubes - as well multiple unconnected cables, which were connected to the room temperature feedthrough but were simply hanging loose inside the 4K chamber. We found that injecting white noise into some of the unconnected cables resulted in an increased axial heating rate. This made us suspect that those connections could be responsible for the anomalously high heating rates observed in setup 1.0 [Leu15].

To avoid this problem in the cryostat 2.0, we now make sure to filter all the cables entering the 4 K chamber<sup>22</sup>. First of all, any unused cables are grounded onto the filterboard. Furthermore, any cables which do not attach to the filterboard - such as the oven or the temperature sensor - are connected through filter feedthroughs on the panel, which simply provide a capacitive path to ground. In Sec. 3.4.1 we list how exactly each connection is filtered and grounded.



**Figure 3.14:** Viewport mounts with slots for heating wire (not shown).

Optical access holes are closed by windows mounted as shown in Figure 3.14.

<sup>20</sup> One panel has a 35 mm diameter hole to allow an ablation beam to enter 12 mm off-axis

<sup>21</sup> 30 mm x 55 mm

<sup>22</sup> Except the Attocube lines

**4K window contamination**

In cryostat 1.0, these windows would periodically fog up. The contamination could be removed by warming up the whole system to the room temperature and back.

To prepare for the contamination, we made the windows heatable. This way, a contaminated window could be individually heated up to remove the fogging. Each viewport sits in an aluminium mount, where a nichrome wire wrapped around the ring allows us to heat up the glass. The metal mount is attached to the chamber panel using Macor to thermally isolate the window from the chamber. In the three years of operation of the cryostat 2.0, we never observed any window contamination, and so it was not necessary to use the heaters. One explanation could be that the KF seals on the outer vacuum chamber were causing the contamination, and so replacing them with CF flanges alleviated the problem. Alternatively, it is also possible that since the windows are no longer well-thermalised to the 4 K chamber, contaminants coat the chamber rather than the viewports.

**3.2.4. Trap mounting**

The trap mounting stack was designed with modularity in mind. The ion trap chip is glued inside a shallow pocket on a carrier printed circuit board (PCB). The electrodes are wire-bonded onto exposed pads of the PCB (Sec. 3.4.1) and the tracks are routed to exposed pads on the back-side. The filterboard features a matching pad pattern. Electrical connection between the two PCBs is established through Fuzz Buttons<sup>23</sup> inserted inside a home-made Macor interposer. Each PCB is screwed onto the interposer, compressing the Fuzz button and establishing a robust connection that is maintained at cryogenic temperatures. The carrier PCB also holds a resistor<sup>24</sup>, which acts both as a trap heater and a temperature sensor. Originally, the trap was surrounded by a copper Faraday cage with small openings for laser access, which was removed when we started testing traps with attached fibres. A fine mesh “roof” shields the trap from the objective and can be used to increase the trap depth (Sec. 3.4.1).

**3.2.5. Atom sources**

In cryostat 2.0, we decided to keep the old cryogenic oven as a backup, while adding two additional loading methods: a resistive oven mounted on a vacuum flange (described above in Sec. 3.1.5) and a target for ablation loading. The resistive oven is a thin-walled stainless steel tube filled with calcium granules. Originally, we intended to eliminate the need for manual oven alignment by placing the tube inside a fixed Macor mount [Pag19]. The electrical connection was made by pressing the oven down with a thin SS flap, and wrapping a copper wire around the fixing screws. While elegant and simple, the oven failed to heat up

<sup>23</sup> Gold-plated Beryllium Copper Fuzz Buttons from Custom Interconnects

<sup>24</sup> PT1000

### 3. EXPERIMENTAL SETUP

efficiently. After iterating on the design, we returned to a more typical construction with a free-floating oven. Doing this, we learned about a few important oven design considerations, which are presented in this section.

#### Oven angle

Theoretically, the neutral flux can come from any direction relative to the photoionization light. However, whenever possible, it is worth orienting the atom flux approximately perpendicular to the neutral fluorescence beam at  $\lambda = 423$  nm in order to minimise the Doppler broadening. Consider atoms of mass  $m$  effusing from an oven at temperature  $T$ . Their speed probability density function is given by [Blu10]

$$f(v) = \frac{m^2}{2k_B^2 T^2} v^3 \exp\left(-\frac{mv^2}{2k_B T}\right),$$

whose mean and standard deviation read

$$\begin{aligned}\langle v \rangle &= \sqrt{\frac{9\pi}{9}} \sqrt{\frac{k_B T}{m}} \approx 1.88 \sqrt{\frac{k_B T}{m}}, \\ \sigma_v &= \sqrt{\frac{32-9\pi}{8}} \sqrt{\frac{k_B T}{m}} \approx 0.68 \sqrt{\frac{k_B T}{m}},\end{aligned}$$

For an angle  $\theta$  between the neutral flux and the laser beam, the resonance is Doppler-broadened by

$$\sigma_f = \frac{\sigma_v \cos(\theta)}{\lambda},$$

where  $\lambda$  is the laser wavelength. When  $\sigma_f$  becomes comparable to the transition natural linewidth  $\sigma_{\text{nat}}$ , the excitation probability decreases.

For an oven at a temperature of  $T = 700$  K effusing  $^{40}\text{Ca}$  atoms, which absorb  $\lambda = 423$  nm light on a transition with natural linewidth  $\sigma_{\text{nat}} = 35.4$  MHz, we find that  $\sigma_f < \sigma_{\text{nat}}$  only for  $\theta = 90 \pm 3$  deg [Luc03]. At an angle of  $\theta = 45$  deg (where we initially mounted an oven), we find that  $\sigma_f \approx 17\sigma_{\text{nat}}$ . Following this observation, we decided to always position the oven approximately normal to the  $\lambda = 423$  nm laser<sup>25</sup>.


#### Oven geometry

To maximise power dissipation and minimise thermal conductivity, a high-resistance oven is beneficial. This is achieved by using a long and thin-walled tube and making it out of a highly resistive material. For a given temperature, the neutral flux increases proportionally to the area of the oven opening (as long as the size of the subsequent aperture is also

<sup>25</sup> The angle with respect to the photoionization laser is irrelevant, as we drive a wide transition to the continuum

increased).

Perhaps a less obvious geometric consideration comes from noting that, for any given current, the oven will be hottest in the centre and coldest at the ends. Therefore, by placing the opening near the end of the tube, we risk that the effective vapour temperature will be reduced. In our designs, we place an opening closer to the middle of the tube, as shown in Fig. 3.15.



**Figure 3.15:** Schematic representation of different oven designs. Thick grey areas represent a metal tube filled with calcium. Thin grey stands for a crimped tube, where the electrical connection is made, and the atom plume is depicted in blue. Top: Oven aperture near the end of the tube. The opening is not as warm as the middle of the tube, which can result in a lowered atom flux. Bottom: Oven aperture near the middle of the tube. The opening is near the warmest region of the tube. The ovens assembled in this thesis are most similar to the middle design.

### Mechanical clamping

To achieve high oven temperatures with low power dissipation, it is necessary to minimise heat conduction to the environment. The dominant heat flow channels are the mounting structure and the electrical cables. In the first design, we placed the oven on a 6 mm-thick Macor spacer, attached directly to the 4 K chamber. Subsequent estimates revealed this to be a terrible idea: despite Macor’s low thermal conductivity, maintaining the oven at  $T = 700$  K would cause a heat flow on the order of 10 W. It should be possible to decrease this number by reducing the contact area, but generally speaking, a rigid cryogenic oven mount with negligible heat conduction would be challenging to build. For this reason, we since suspended the cryogenic ovens only by their cables. The room-temperature oven was initially positioned on a Macor mount, but we removed it as well, which allowed us to operate it at a lower current. To provide mechanical stability, we instead clamp the oven wires (see Fig. 3.16).

### Electrical connection


In the initial design, the oven tube was directly connected to cryogenic copper cables. While this created a good electrical connection, the thermal connection was too good.



### 3. EXPERIMENTAL SETUP

Later calculations revealed that, since the oven cables were thermalised to 4 K, this design effectively held the oven ends at cryogenic temperatures. To address this problem, we retreated to a more traditional design, where an intermediate wire (made out of tantalum) is introduced between the oven and the cables to sustain large temperature gradients. The exact optimal length of the tantalum wires can be, in principle, calculated. However, the model depends on the exact boundary conditions introduced by copper cables and gets furthermore complicated by the temperature dependence of thermal and electrical properties of copper, tantalum and SS. Simplified calculations<sup>26</sup> indicated that, for our set of parameters, the tantalum wire should be of comparable length to the steel tube, and that increasing either length generally increases the temperature in the middle of the oven for a fixed power. Consequently, one can think of the tantalum wire as a flexible way to make the oven longer. In experiments, we generally make the tantalum length comparable to the oven tube length.

#### Employed designs



**Figure 3.16:** Atomic sources inside the 4 K chamber. Neutral plumes travel approximately perpendicularly to the  $\lambda = 423$  nm neutral fluorescence laser. In the image, the assembly is not finished: the resistive is later moved to  $\approx 2$  mm to the side, in order to clear the optical axis.

With those considerations in mind, we employed the atom sources as shown in Fig. 3.16.

<sup>26</sup> Assuming copper wires are held at a fixed temperature, and all the material parameters keep their room-temperature values



The resistive oven is a  $\approx 20$  mm-long SS tube with outer diameter of 1 mm and 100  $\mu\text{m}$  thickness<sup>27</sup>. The tube is opened and bent close to its centre, and its ends are crimped. Tantalum wires<sup>28</sup> of  $\approx 20$  mm length are spot-welded to each end of the tube and clamped to rigid insulated copper cables, which are then clamped through a Macor piece to provide mechanical stability. The oven aperture is placed  $\approx 30$  mm away from the trap chip centre. An additional aperture of 10 mm height is placed 20 mm away from the chip centre to collimate the atom beam. The oven is positioned by hand by bending the copper and tantalum wires. The tube is placed  $\approx 2$  mm off the optical axis to allow optical access through the chip centre. This results in  $\theta \approx 86$  deg angle between the neutral flux and the neutral fluorescence beam, producing minimal Doppler broadening.

The ablation target is placed on the other side of the trap chip. The block is held with a screw, its flat side oriented towards the centre. The ablation beam travels along the long side of the trap and focuses on the target. The plume of atoms then leaves the target perpendicular to the surface and flies across the chip at  $\theta \approx 84.5$  deg angle to the neutral fluorescence beam. However, since the ablated atoms travel at speeds much higher than typical for thermal sources, even this small angle results in a clear Doppler shift (Sec. 5.1.3). A Faraday cage hole (not shown in in Fig. 3.16) was used to aperture the ablation plume.

### 3.3. Ion traps

As shown in Fig. 3.1, we worked with multiple different traps throughout this thesis. While I did not design or fabricate any of the traps myself, this section summarises their main features.

#### 3.3.1. Gold on quartz trap

The only trap used in the setup 1.0 was a gold-on-quartz chip shown in Fig 3.17. The trap was based on a 5-wire asymmetric design [Chi05], with the outer DCs segmented into nine electrodes each. The gold electrodes were 8  $\mu\text{m}$  thick with 6  $\mu\text{m}$  gaps, and the ion was trapped 50  $\mu\text{m}$  above the surface. The design was intended for a single-zone operation close to the centre, with additional control electrodes used for fast potential modulation. The whole chip was 10 mm  $\times$  10 mm in size. Laser access was possible along the trap axis and at  $\pm 45$  deg, while the direction perpendicular to the trap axis was blocked by wirebonds. The trap was designed by F. Leupold and J. Alonso, and fabricated by J. Alonso in the ETH Zurich cleanroom [Leu15].


Unlike the later traps, which can be easily swapped, the original trap was pretty much fixed in place by the electrical connections. The wirebonds connected the trap to larger gold wires, which were mechanically fixed by a PEEK interposer and soldered onto the original filterboard.

---

<sup>27</sup> McMaster-Carr 89935K63

<sup>28</sup> 0.5 mm diameter

### 3. EXPERIMENTAL SETUP



**Figure 3.17:** Gold on quartz trap in the setup 1.0. a) Picture of the trap wirebonded onto the interposer and soldered onto the filterboard b) Trap electrode layout. Electrodes shown in orange and green supply DC voltages, and the yellow electrode is the RF.

The quality of the coherent operations achieved in this trap is discussed at length in [Leu15], and Chapter 7 presents the non-classicality experiments which were also done in this original system. Aside from the technical noise, which we believe originated elsewhere in the setup, the trap itself functioned very well and we did not observe any laser-induced charging. Unfortunately, the fast potential switching experiments – for which the system was originally designed – hit some fundamental roadblocks, which are discussed in the thesis of C. Zhang [Zha21]. Eventually, we scrapped this project and upgraded the setup to operate traps with more novel features.

#### 3.3.2. CMOS trap

The first trap we attempted to use after the experiment upgrade was designed by C. Zhang during his Master’s thesis. The chip was our exploration into the space of fabricating ion traps with commercial CMOS processes [Meh14]. The trap implemented a symmetric 5-wire design, with 31 DC electrodes and the RF electrodes shaped to minimise the axial pseudopotential curvature.

We experienced a lot of problems with reliable wire-bonding of the chip to the carrier PCB. Initially, we aimed to avoid using a magnetic nickel layer common to many standard PCB finishes. We therefore used an immersion silver finish, which should adhere to aluminium wirebonds. However, the process proved difficult to optimise, with at least 10% of wirebonds unconnected. Owing to the small chip and pad dimensions, we also couldn’t place multiple bonds per pad. Subsequently, we replaced the carrier PCB finish with EPENIG<sup>29</sup>, which allowed for reliable wirebonding.


After the wirebonding problems were resolved, an outstanding issue was a non-Lorentzian lineshape of the resonator reflection curve at temperatures below 20 K. We speculate that

<sup>29</sup> Electroless Nickel Electroless Palladium Immersion Gold

this was caused by a design error: some regions of the trap did not include a ground plane to shield the RF signals from the silicon substrate. We attempted to trap for several weeks, both at 4 K and with the chip locally heated to above 20 K, but never observed any signatures of ions. Around that time, we received the first samples of the integrated optics traps, and we decided to try trapping in those instead.

### 3.3.3. Waveguide-integrated traps

The traps with integrated waveguides were designed by K. Mehta and fabricated by a commercial foundry<sup>30</sup>. In this section, I describe the electrical and mechanical design, while Sec. 4.6 goes into the details of the waveguide components. The electrode layout, shown in Fig. 3.18, implements a symmetric six-wire design with three zones [Chi05]. The outer DC electrodes are segmented into nine electrodes each, and the centre DC electrode is split in two. At each zone, two large openings in the centre DC electrodes expose the ion to the light from the chip. The stack cross-section is shown in Fig. 3.18. The gold (Au) electrodes



**Figure 3.18:** (left) Microscope image of the waveguide-integrated ions trap wirebonded onto the carrier PCB and connected to the fibre array. (right) Cross-sectional cut of the layer stack. Note that the Si substrate extends significantly below the image, and forms the majority of the chip.

are only 300 nm thick, and are deposited on a glass ( $\text{SiO}_2$ ) dielectric. A platinum (Pt) ground plane shields the electrodes from the underlying silicon (Si), although the ground plane has openings to allow grating light to come out. Silicon-nitride ( $\text{Si}_3\text{N}_4$ ) waveguides are patterned in two layers underneath the ground plane, and the structure is fabricated on a silicon substrate.

Note that there is no electrical shielding of the grating coupler, and the ion has line-of-sight exposure to the glass, the waveguides, and the substrate. Additionally, along the whole trap axis, the dielectric of the gap between the electrodes is directly below the ion. This can influence the heating rates and cause charging [Har10; Tel21]. Originally, we were really worried that this dielectric exposure, which amounts to  $\approx 20\%$  of the solid angle, would prevent us from trapping at all. This turned out not to be the case, but substrate charging remains a major limitation and annoyance for the experiments (Sec. 8.2.2).

<sup>30</sup> LioniX International

### 3. EXPERIMENTAL SETUP


The first two traps we tested had the same the top metal layer as used in the final design but did not include the  $\text{Si}_3\text{N}_4$  waveguides. This way, we had time to figure out the ion trapping without worrying about unnecessary complications, while simultaneously working on the optimisation of the fibre-waveguide coupling.

#### Trap #1: unsuccessful trapping

We spent about three months, and multiple cool-downs, trying to trap in the trap #1. Initially, we found it challenging to be confident about the beam alignment since the trap was very reflective, which made it difficult to see diffuse scatter. Sec. 4.5.3 describes the laser and imaging alignment method we settled on. In the end, we discovered an accidental large focal shift of the  $\lambda = 375$  nm laser, which meant that while other lasers were focused onto the trap location, this laser was scattering off the chip, thus charging the trap surface over many months (Sec. 4.5). We then decided to swap traps before trying to trap again.

#### Trap #2: ion trapped!

In addition to reducing  $\lambda = 375$  nm scatter, we implemented two changes. First, we placed four drops of conductive epoxy at the four corners of the trap, to better ground the Si substrate. Second, we placed the resonator PCB in a copper box to reduce the amount of RF radiation picked up by the DC electrodes. One of these three changes allowed us to finally start trapping ions.




**Figure 3.19:** Stray curvature found near the metal layer opening. (Red) In Trap #2 the stray potential is positive (confining). (Blue) In subsequent traps, the stray potential is negative (anti-confining). The measured anti-confinement is comparable for traps #3 to #5.

We only performed a rudimentary set of measurements on Trap #2 in order to compare its performance to the setup with integrated light. The main surprise was a strong positive stray axial potential curvature around the coupler openings (Fig. 3.19). This resulted in an

axial centre-of-mass frequency  $\omega = 2\pi \times 2.06$  MHz even when all the DC electrodes were grounded! We also performed a set of heating rate measurements (Sec. 5.6.4) for comparison with traps with integrated waveguides.

### Trap #3: Rabi flopping through integrated waveguides

The next trap included fibres and integrated waveguides. Shortly after the first cool-down, we discovered an issue with the fibre feedthrough and decided to warm up the cryostat and open the chamber. During the warm-up, two of the electrodes got damaged as shown in Fig. 3.20. While scanning the literature for possible explanations, we stumbled onto a



**Figure 3.20:** Metal layer damage following a warm-up and chamber opening with voltages applied to electrodes.

suggestion that such an effect could arise due to a phenomenon of *electrochemical migration* caused by water contamination (from humid air) in combination with a voltage applied to electrodes [Dom08]. Indeed, we did set the DC voltage sources to 0 V before exposing the chamber to air. After confirming that no electrical shorts occurred, we nonetheless continued to operate the trap but working in a zone which was not adjacent to the damaged electrodes.

After trapping, we were able to perform the first coherent Rabi oscillations using light delivered through the waveguides. We found that the substrate on this trap and all subsequent traps still produced stray curvature, but this time it was anti-confining (Fig. 3.19). We also found and mitigated charge transients caused by silicon photoconductivity (Sec. 8.2.1).

We then made the first attempts at loading the trap from an ablation source, which are discussed in Sec. 5.1.3. While we were able to load ions using this technique, it was not efficient and lead to significant substrate charging.

### Trap #4: Two-qubit gates

The next trap was a clean sample without damaged electrodes or ablation-induced charging. Shortly after installing the trap, we were able to tune up and benchmark our quantum operations more thoroughly, as described in Chapter 5. A few months later we succeeded at

### 3. EXPERIMENTAL SETUP

creating the first entangled states using  $\lambda = 729$  nm delivered through integrated waveguides, as presented in Chapter 8.

#### Trap #5: Dissipative entanglement

After operating trap #4 for several months, the ion lifetime started to deteriorate, and we decided to fix several suspected issues. Since the fibres attached to the trap were epoxied through a vacuum flange, the trap had to be discarded as well. After installing a room-temperature oven and a reusable fibre feedthrough (Sec. 4.6.4), we cooled down the final trap #5. This chip stayed in vacuum for a year, and was used in the dissipative entanglement experiments described in Chapter 9. This trap behaved very similarly to the trap #4, except for a puzzling discrepancy in the effective anharmonicity of the trap potential (Sec. 9.3.3).

## 3.4. Electrical control

In this section, I describe all the electronics used to control our experiments. I will first describe the components closest to the ion trap, and then gradually zoom out to see the whole picture. Sec. 3.4.1 describes the electronics in the 4K chamber. The following Sec. 3.4.2 describes the devices outside the vacuum system, and how these are connected to the inside.

### 3.4.1. Cryogenic electronics

Fig. 3.21 presents the schematic of electronic connections inside the 4 K chamber. All the connections to the ion trap are made through a single filterboard which routes the DC, RF, Zeeman and trap heater lines. The soldered filterboard is shown in Fig. 3.22.

#### DC connections

##### Original DC connectors

In the original design, the DC connections were delivered to the filterboard through card-edge connectors. We found it very challenging to reliably slot them in and out. The filter outputs were soldered and wirebonded directly to the trap, so swapping traps meant replacing the filterboard as well.

In the upgraded design, the DC voltages are delivered with a micro-D 51-pin cable<sup>31</sup> onto the filterboard. Twenty-nine voltage lines are cleaned with a first-order “low-pass” RC filter, while the rest are grounded. The experiments presented in this thesis used  $R = 180\ \Omega$  and  $C = 33\ \text{nF}$  to achieve a cut-off frequency of  $\approx (2\pi) \times 27\ \text{kHz}$ . The capacitors are all

<sup>31</sup> Home-made cables with connectors from Glenair

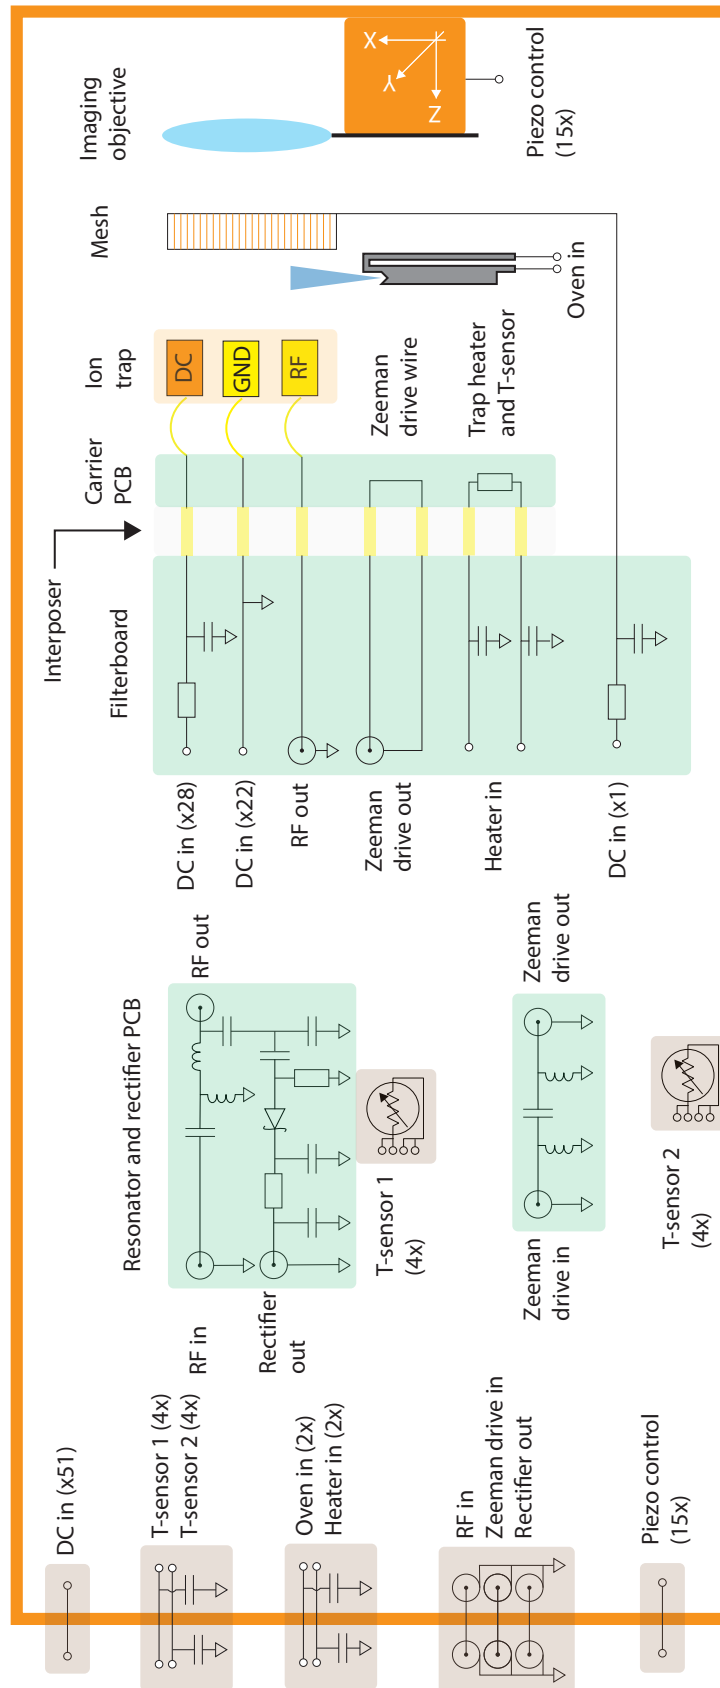



Figure 3.21: Overview of cryogenic electronics

### 3. EXPERIMENTAL SETUP



**Figure 3.22:** (Left) Trap filterboard PCB (ion trap placed on the back side) (Right) Trap assembly right before it is inserted inside the 4K chamber. The trap is shielded from an imaging objective (not shown) is a transparent gold mesh. The mesh is connected to a dedicated DC line.

C0G (NP0) ceramic, which has the most stable performance at cryogenic temperatures [Tey10]. All the signal lines and a ground line are connected through an interposer to the carrier PCB and wire-bonded onto the trap. The ground connection is made through four wire-bonds in the four corners of the trap, and additional conducting epoxy grounds the silicon trap substrate [Meh20].

One DC connection does not go to the trap but is instead routed to a mesh which shields the ion from the imaging objective. Originally, the mesh was placed 2 mm above the chip, and simulations indicated a 6-fold increase in trap depth when +10 V was applied to it [Leu15]. We therefore used the mesh as an electrode to increase the loading rate. Once we started using traps with integrated optics, the mesh had to be moved up to make room for the fibre attach. Now it is placed 10 mm above the chip and has little influence over the trap depth.

#### RF circuit

##### Helical resonator

The old RF resonator was a bulky helical resonator based on a design in [Siv11]. It had a high quality factor, but was difficult to impedance-match and took up a lot of space. It was also designed for dual-species  $^{40}\text{Ca}^+$ - $^9\text{Be}^+$  operation, and thus its resonance frequency was much higher than necessary for trapping  $^{40}\text{Ca}^+$  alone.

Strength: High quality factor.



High electrical conductivity of copper at cryogenic temperatures allowed to increase the resonator  $Q$ -value, and hence the voltage gain. We obtained  $Q \approx 400$  at  $2\pi \times 105$  MHz, as measured at 4 K with a test 6 pF capacitor. However, a reduced value of  $Q \approx 100$  at  $2\pi \times 93$  MHz was obtained with the trap connected and the resonator placed inside the setup. This happened largely because the trap and the filter board did not create a purely capacitive load.

#### Weakness: impedance matching

The helical resonator was matched to a  $50\ \Omega$  source by translating and rotating a small in-coupler antenna while measuring reflection with a spectrum analyser. Despite improving the in-coupler's mechanical stability by increasing coil thickness and locking it with a set screw, the impedance matching was unreliable since the antenna could only be adjusted when the 4 K chamber was open, and subsequent cooldown always introduced significant shifts. The best impedance matching at cryogenic temperatures was obtained by introducing a small deliberate mismatch at room temperature. However, it was difficult to make this process reproducible.

The upgraded resonator is a small circuit board. The resonator was built by C. Zhang based on a design presented in [Bra16c], and more details can be found in his thesis [Zha21]. RF voltage at  $\omega = 2\pi \times 34$  MHz is delivered with an SMA cable and amplified by a factor of  $g \approx 14$  with a PCB resonator. We used Rogers as a PCB substrate to minimise RF losses. Impedance matching is achieved over a few temperature cycles by measuring the change of the reflection coefficient at cold temperatures, modelling the change in system capacitance and inductance, and adjusting the values of the matching elements accordingly. The resonator PCB is enclosed in a grounded copper housing<sup>32</sup>, and its output is soldered onto the filterboard through a short wire. On the filterboard, we removed the ground plane underneath and around the RF track to minimise load capacitance.

To monitor and calibrate the applied RF voltage, the resonator output is picked off with a capacitive voltage divider and rectified to DC with a Zener diode [Joh16; Reu17]. This monitor signal leaves the resonator chip with a separate SMA cable. While we found systematic differences between this monitor and the RF voltage inferred from the trap frequency measurements, it is nonetheless a very useful tool for approximately setting the trap voltage and monitoring any drifts. We have not yet attempted to use the rectifier for active trap potential stabilisation<sup>33</sup>.


<sup>32</sup> We introduced the housing after a few months of unsuccessful trapping attempt to reduce the RF pickup on DC electrodes

<sup>33</sup> If I were to do that again, I would pick off RF voltage and only rectify it close to the controller to avoid low-frequency noise pickup

### 3. EXPERIMENTAL SETUP

#### Zeeman drive

We employ a current-carrying wire to produce an oscillating magnetic field that directly drives the Zeeman qubit. To that end, a signal at the Zeeman qubit frequency is delivered with an SMA cable and filtered with a band-pass current filter, again described in the thesis of C. Zhang [Zha21]. The output of the Zeeman filter PCB is screwed onto the filterboard with an SMA elbow. The oscillating current then passes through a track located on the carrier PCB underneath the ion trap chip, as shown in Fig. 3.23.



**Figure 3.23:** a) Side view of the current track in the trap carrier PCB. The oscillating current creates a vertically oscillating B-field at the location of the ion 1.1 mm away. b) Top view of the trap wire-bonded into the carrier PCB. The red line shows the current track underneath the ion trap.

The two-wire configuration creates a magnetic field oscillating perpendicular to the trap surface. Since the quantization axis is aligned with the chip, this configuration drives the Zeeman qubit  $\Delta m_J = 1$  transition efficiently. The ground plane of the trap provides shielding from electric field noise caused by finite track impedance.

We can calculate the coupling strength by assuming the wires to be infinitely long. At a B-field of 5.9 G the Zeeman splitting is  $2\pi \times 16.5$  MHz. This corresponds to a wavelength of 20 m, and so we can consider the currents in the two tracks to be in phase to a good level of approximation. This current creates an oscillating field of  $\approx 3$  G/A, resulting in a Rabi frequency of:

$$\Omega = 2\pi \times 4.3 \text{ kHz/mA}.$$

We typically operate at a Rabi frequency  $\Omega = 2\pi \times 76$  kHz, which corresponds to  $\approx 18$  mA on the track.

#### Trap heater

The back-side of the trap carrier PCB houses a PT1000 resistor that we use as a heater and a local temperature sensor. Since the interposer is a thermal insulator, this allows us to locally raise the trap temperature without affecting the vacuum environment. We used this method to measure the trap temperature dependence of heating rates from 4K up

to 150K (Sec. 5.6.4). We also regularly use the heater to keep the trap warmer than its surrounding during cool-downs, which should reduce cryo-pumping and contamination of the chip surface. The temperature can be measured by reading out the heater resistance, although calibration curves are only available down to 77 K.

To avoid putting large DC currents through the system ground path, the heater current is delivered through a floating twisted cable pair that plugs into the filterboard with pin heaters. Additional  $C = 100$  nF shunt capacitors provide a ground path for high-frequency noise and RF pick-up.

### Other electronics

Other connections are routed outside the filterboard. A twisted pair delivers current to the resistive oven. Two temperature sensors are placed inside the chamber – one on the base and one on the resonator – with their leads attached to a single D-sub 9-pin connector. All of these connections are filtered on the chamber panel feedthroughs. Finally, the imaging objective nano-positioner x/y/z control and read-out lines are connected to a single D-sub 15-pin connector. We were advised by the manufacturer<sup>34</sup> against shunting those, as it could cause problems in movement or position readout. We do however ground the body of the nano-positioner by directly screwing it into the chamber.

#### 3.4.2. Signal delivery

An overview of main signal delivery channels is shown in Fig. 3.24.

### Grounding considerations

We opted for a “star ground” configuration, where the 4K chamber (which is an excellent conductor) serves as the voltage reference. Unfortunately, the 4K chamber is unavoidably connected to the building earth through the helium dewar and the PT. We therefore embraced the fact that the vacuum chamber is connected to earth, and did not attempt to isolate it from the (earthed) optical table.

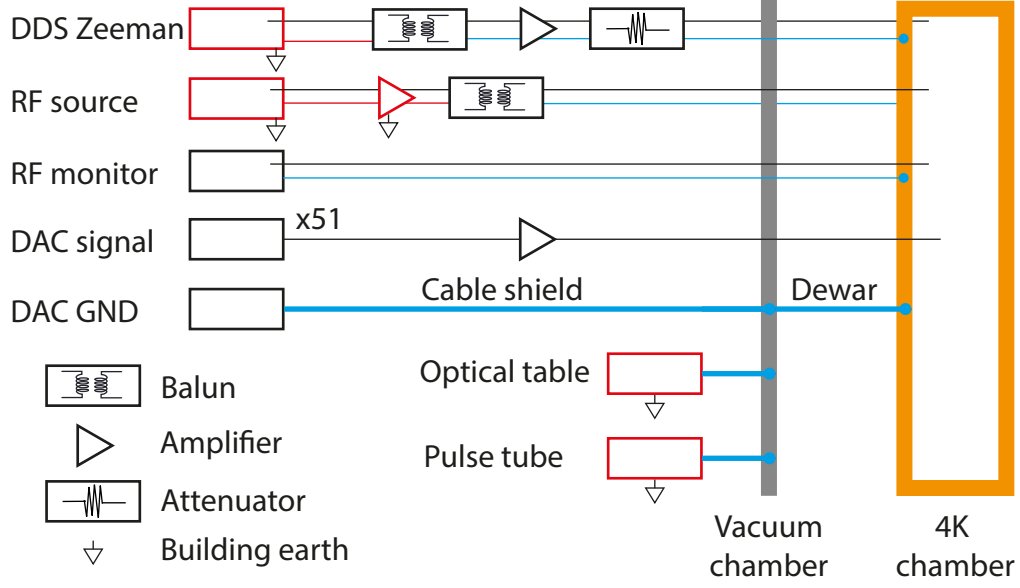
All signals are sent differentially, with the return current path connected to the ground only at the 4K chamber. Whenever a source forces an earth connection – as is the case for the RF synthesiser, the RF amplifier and the DDSs – we place a balanced-to-unbalanced transformer (balun) in the chain to avoid ground loops. These devices are now placed all around the lab in a way shown in Fig. 3.25.

The DC current return path is created by a single low-resistance connection from the 4K chamber to the DAC. The helium dewar is used as the lowest-resistance line between the 4K chamber and the vacuum chamber. We then use a thick metal braid to connect the vacuum chamber to the DAC (and to shield the DC cable simultaneously). Unused DAC lines can also be optionally used as ground connections.

---

<sup>34</sup> attocube systems AG

### 3. EXPERIMENTAL SETUP



**Figure 3.24:** Cryogenic signal delivery overview. Blue lines indicate ground connections, with the lowest-resistance ground connections (DC current return path) highlighted in thick blue. Red elements indicate connections to the building earth.

#### Cryogenic cabling considerations


When connecting a cryogenic environment to a room temperature environment, one must pay attention to avoid excessive heat load. Since good electrical conductors are usually good thermal conductors as well, most standard cables conduct away a lot of heat. This is not a big problem for a cryostat with a lot of cooling power and a few cables but can become an issue when scaling to larger numbers of connections [Kri18].

Consider a cable of length  $d$  and cross-sectional area  $A$  connecting a hot reservoir of temperature  $T_H$  with a cold reservoir of temperature  $T_C$ . From the Fourier equation, the amount of heat flowing into the cold end per unit time:

$$\dot{Q} = \frac{A}{d} \int_{T_C}^{T_H} \kappa(T) dT,$$

where  $\kappa(T)$  is the temperature-dependent thermal conductivity of the cable. The quantity  $\int_{T_C}^{T_H} \kappa(T) dT$  is known as the integrated thermal conductivity of a material and can be looked up in engineering tables [Dut15]. The expression for the total DC resistance  $R$  is more complicated since the cable thermalises with a conductivity-dependent equilibrium temperature  $T(x)$  at point  $0 < x < d$  and the total resistance requires integrating over temperature-dependent conductivity.

To gain intuition about a general material, we assume that the thermal conductivity



**Figure 3.25:** A balun during the setup process. The unbalanced end (left) is connected to the body of the transformer, which may or may not be in contact with the earth ground. After connecting a coaxial cable to the floating (balanced) output (right), we clamp a plastic hose around it. This way, the balun can be moved around without the balanced end contacting any metal pieces and shorting to ground.

decreases with temperature. This allows us to write:

$$\dot{Q} < \frac{A}{d} \Delta T \kappa(T_H),$$

where  $\Delta T = T_H - T_C$ . Likewise, the electrical resistivity generally decreases with temperature, leading to

$$R < \gamma \frac{d}{A} \rho(T_H),$$

where  $\rho(T_H)$  is the electrical resistivity, which can be related to  $\kappa$  through the Wiedemann-Franz law:

$$\kappa(T) \rho(T) = LT$$


where  $L = 2.44 \times 10^{-8} \text{ W } \Omega / \text{K}^2$ . While the Wiedemann-Franz law is not exact, it is a decent approximation for most metals at most temperatures [Sim13]. Combining these expressions gives:

$$\dot{Q} < L \frac{\Delta T T_H}{R}. \quad (3.1)$$

which allows us to bound the heat loss for a given DC resistance<sup>35</sup>. This is shown in Fig 3.26.

<sup>35</sup> Note that this  $R$  is the resistance at the equilibrium temperature, and will be lower than the resistance

### 3. EXPERIMENTAL SETUP



**Figure 3.26:** Upper bound on heat flow caused by a cable of resistance  $R$ , obtained by evaluating the right-hand-side of the inequality in Eq. (3.1).

This calculation reveals several insights. First of all, we can minimise the heat load on the 4K stage by thermalising the cables to the 1st stage heat exchanger, assumed to be at 55 K. This allows us to install a few low-resistance connections since the 1st stage can handle significantly larger loads. Installing a standard low-resistance cable (with 1–10 mΩ resistance) directly between the cryogenic stage  $T_C = 4$  K and the room-temperature environment ( $T_H = 300$  K) is a bad idea, as it can conduct  $> 100$  mW of power. On the other hand, 100 wires with 10 Ω resistance can result in a very acceptable heat load of  $< 1$  mW even when not thermalised at the first stage. We follow those design principles in our experimental layout, and all lines are clamped at the 1st stage exchanger before travelling down to the 4K chamber.

#### DC lines

Up to 51 DC voltages are generated by a homemade high-precision voltage source<sup>36</sup> and controlled through a serial port. The voltage source is based on a 20-bit AD5791 DAC, with  $\pm 10$  V output and  $< 0.1$  nV/Hz<sup>1/2</sup> measured noise at  $\omega = 2\pi \times 0.5$  MHz (post-filter). The output voltage was sufficient to produce the desired  $\omega = 2\pi \times 2$  MHz axial confinement in most traps but was only marginally sufficient for the integrated optics traps with strong stray anti-confinement. We then added an intermediate amplifier between the voltage source and the trap. It features 51 channels of ADA4522 op-amps in a non-inverting configuration with a gain of 2.5 and a post-filter at 8 kHz to remove an op-amp noise peak at 800 kHz. Pin headers connect the channels which are in use to output cables, and the unused channels are left with their outputs floating. The DC signals are then delivered through two cables to a pair of Micro-D 25-pin vacuum feedthroughs, with the ground connected through the cable

at  $T_H$

<sup>36</sup> Built by C. Zhang following a design by P. Maerki

shield to the vacuum chamber body. Inside the vacuum, the DC lines are delivered through high-resistance Manganin wires of  $d = 1.4$  m,  $R = 70\ \Omega$ . The wires are taped together with Kapton tape and thermalised at the 1st stage heat exchanger by a copper clamp. The wires are soldered to a single Micro-D 51-pin connector on the 4K chamber side, which plugs into the 4K chamber feedthrough.

### Coaxial lines

There are three coaxial lines (RF, rectifier, and the Zeeman drive) connected to the 4K chamber. The RF signal is generated by a low-noise signal synthesiser. After an optional “tickling” circuitry [Leu15], the signal is passed through a balun to remove a ground loop and sent to the vacuum system with a coaxial cable. The RF monitor signal exits the vacuum system with a coaxial cable and is plugged into a (floating) voltage logging device. The Zeeman drive pulses at  $\omega = 2\pi \times 16.5$  MHz are synthesised by mixing two phase-coherent DDS sources at  $\omega_1 = 2\pi \times 150$  MHz and  $\omega_2 = \omega_1 + \omega$ . The signal is then low-pass-filtered to remove the component at  $\omega_1 + \omega_2$ , passed through a balun and sent to the vacuum chamber.

All the three lines enter the vacuum chamber through a ground-isolating feedthrough and connect to vacuum coaxial cables.

#### Cryo-coaxial

The coaxial cables used in the setup 1.0 were discovered to be  $75\ \Omega$  models and subsequently discarded. No wonder impedance matching was never great in the setup 1.0!

The replacement high-resistance semi-rigid cables were problematic since their shields were not insulated. To avoid shorts and ground loops, we taped the cables with Kapton, but this was not reliable and made us worry about virtual leaks from the tape. In the current iteration, we took a hybrid approach. The cables are all low-resistance UHV-compatible insulated coaxial cables. After thermalising at 55 K, the coaxial lines are attached to the 4K stage through high-resistance semi-rigid coaxial elements, which reduces the heat load on the 4K stage.

### Current-carrying lines

The oven and heater lines are driven by remote-controlled power supplies. The cables are connected through a single filtered D-sub 8-pin feedthrough. Those wires are an exception to the design considerations presented earlier in this section; it is counter-productive to increase the resistance of lines that carry significant currents since Joule heating can easily overwhelm thermal conduction. Therefore, our oven and heater lines are simply twisted pairs of 0.5 mm diameter copper wires ( $R \approx 30$  m $\Omega$ ), terminated with pin headers.

### 3. EXPERIMENTAL SETUP

#### Temperature sensors

There are a total of four temperature sensors in the cryogenic system: one is placed on the 2nd stage exchanger inside the PT, and three in the vacuum chamber. After the setup upgrade, we placed them in various positions to verify system performance. Nowadays, we keep one in the 4K chamber, one on top of the chamber and one on the 1st stage heat exchanger. The temperature sensors are all calibrated silicon diodes<sup>37</sup>, whose forward voltage drop is proportional to the temperature. Aside from the thermal load considerations, the cables connecting to the sensors should be resistive to allow the sensor to thermalise to its environment – we learned this the hard way. Thin copper wires offer a good compromise by minimising heat flow while introducing little added resistance that can bias the temperature readout. Our temperature sensors are connected through thin 0.1 mm diameter copper cables with about  $2\,\Omega$  resistance<sup>38</sup>.

#### Piezo control

The imaging objective is mounted on a 3-axis piezo stack, whose position can be adjusted and capacitively measured with a dedicated controller. The drive and measurement signals - five per axis - enter vacuum via a separate Micro-D 15-pin connector, and a 15-wire cable provided by the manufacturer takes the signals down to the base of the 4K chamber. This is the only connection we have not modified. However, during a recent cool-down, one of the cables broke, and we now lack the ability to read out one axis<sup>39</sup>. The cable is also stiff and difficult to handle in confined spaces, so I would recommend making a new one.

#### 3.4.3. Magnetic field coils and rings

Beginning around the year 2016, many groups in the ion trap community began using permanent magnets produce magnetic fields, after it was shown that their noise can be similar or lower than the noise of current coils [Rus16]. Since our experiments do not require fast modulation of the B-field strength, we also decided to trade the flexible B-field control for superior coherence. Our permanent magnet “coil” is a circular array of small cylindrical magnets. We model it as a set of  $N$  magnetic dipoles arranged at positions  $\vec{R}_i$ . A dipole with magnetic moment  $\vec{m}$  at position  $\vec{R}_i$  creates a magnetic scalar potential:

$$\psi_i = \frac{\vec{m} \cdot (\vec{R}_i - \vec{R}_0)}{|\vec{R}_i - \vec{R}_0|^2}.$$

---

<sup>37</sup> Si-410 from Scientific Instruments

<sup>38</sup> In hindsight, since we decided to read out the voltage through a four-wire measurement, the cable resistance should not introduce any bias, and lower-conductivity manganese wires can be used instead.

<sup>39</sup> It is not the first time we lost the position read-out, and we are used to working like this by now



at position  $\vec{R}_0$  [Gri17]. The total field is calculated from:

$$\psi = \sum_{i=1}^N \psi_i \quad \text{and} \quad \vec{B} = \mu_0 \vec{\nabla} \psi.$$


The magnetic moment of a small permanent magnet with volume  $V$  magnetised along  $\hat{z}$  is given by:

$$\vec{m} = V B_r \hat{z} / \mu_0,$$

where  $B_r$  is the magnetic remanence of the material. Consider a pair of permanent magnet coils placed at  $\pm \hat{z}d$  from the trap. Each coil consists of  $N$  permanent magnets with magnetisation  $\vec{m}$ , arranged in a circle of radius  $r = \eta d$ . One can now calculate that the total field at the centre is given by  $B = B_0 \hat{z}$ , where

$$B_0 = \frac{N B_r V}{d^3} \frac{2 - \eta^2}{2\pi(1 + \eta^2)^{5/2}}.$$

At the exact centre of the coils, the B-field gradient is zero. However, to ensure robustness to misalignment, the ratio  $\eta$  can be selected such as to minimise the curvature of the B-field, thus minimising the gradient in a large region of space. Setting  $\frac{\partial^2 B}{\partial z^2} = 0$  reveals a ‘‘Helmholtz-like’’ configuration can be achieved for  $\eta = 0.59$  or  $\eta = 2.76$ , and the first solution is generally preferred as it produces a stronger field. Interestingly, the permanent magnet coils can be smaller than the traditional Helmholtz coil, which requires  $\eta = 2$ . Fig. 3.27 illustrates the B-field orientations for different values of  $\eta$ .



**Figure 3.27:** B-fields created by permanent magnet coils located at  $\pm d = 142 \text{ mm}$  for increasing values of  $\eta$ . B-field direction is indicated by the blue lines, and the magnitude by the colour, with white corresponding to the largest fields.  $\eta = 0.84$  produces a large region of B-field aligned primarily along  $\hat{z}$ . Is it also the configuration with the lowest field curvature along  $\hat{z}$ , but that is not clearly visible from the plots.

### 3. EXPERIMENTAL SETUP

The original design used a pair of coils with ion-coil spacing of  $d = 142$  mm and coil radius  $r = 0.59d \approx 84$  mm. We opted for SmCo magnets, which have a lowest temperature sensitivity of any permanent magnet material ( $\Delta B_r/B_r \approx 5 \times 10^{-4}/\text{K}$ ). Each magnet<sup>40</sup> is a cylinder of 3 mm radius and 4 mm height, and an average specified remanence of  $B_r = 1$  T. We used  $N = 52$  magnets per coil to create a field of  $B_0 = 4.25$  G. Each coil is an aluminium ring with machined cylindrical slots to which the magnets are glued.

During the cryostat upgrade, we had to increase the ion-coil distance, and we simultaneously decided to increase the magnetic field. We now place  $N = 192$  magnets per coil in two circles with average radius  $r = 103$  mm, producing  $B_0 = 5.9$  G at  $d \approx 150$  mm. The ratio  $\eta = r/d = 0.68$  is no longer optimal for curvature suppression. We deliberately decided to ignore this design constraint, since the trap axis is oriented at 45 deg to the quantization axis, and no value of  $\eta$  allows to eliminate curvature along that direction. In the experiments, we record a magnetic field gradient of  $\approx 3$  G/m along the trap axis (Sec. 9.3.4). This gradient can be reproduced in simulations by introducing  $\approx 5$  mm of misalignment between the coils and the trap centre, which is consistent with the observed trap-chamber misalignment and the coil positioning accuracy. We also observe long-term B-field drifts consistent with the specified thermal dependence of  $B_r$ , and thus we think that magnetic field stability can be improved by temperature-controlling the coils.

#### 3.4.4. Device control

To operate the experiments, a large number of electrical and optical devices need to work at the same time. Some require dynamical control with well-defined timing, while other tasks can be handled asynchronously.

##### Synchronous control

Experimentally time-critical components are synchronised and controlled by a master FPGA board programmed from a control PC [Neg18b]. It allows generating the input and output signals at well-defined intervals. In our experiments, synchronous i/o was used for:

- Generating laser pulses
- Generating Zeeman drive pulses
- Modulating RF for micromotion compensation (a.k.a. “RF tickling”)
- Generating TTL signals for pulse sequence monitoring and occasionally for device triggering (e.g. Sec. 5.1.3)
- Reading out the ion fluorescence signal from the PMT
- Real-time measurement randomisation (Chapter 7)

---

<sup>40</sup> DeltaMagnet from IBS Magnet

The output pulses are generated phase-synchronously by DDS modules, controlled by the FPGA and clocked with an external atomic clock [Neg18b]. We have twelve available channels across three cards. Each channel can only output one frequency, and so bichromatic control requires two channels to be combined. Low-frequency signals likewise cannot be directly generated, so RF tickling and Zeeman drives require mixing multiple channels. The maximum output power of a DDS source is 0 dBm, and so the DDS output is usually followed by an amplifier.

One important consideration is that if two DDS channels are triggered simultaneously, the relevant laser pulses will likely not be applied at the same time. This happens due to other differential delays in the chain, e.g. different cable lengths ( $\approx 5$  ns per metre of cable length difference) or acousto-optic modulator (AOM) rise times ( $\approx 100$  ns for 0.5 mm distance difference from the transducer, Sec. 4.1.3). In our experiments, it was not necessary to calibrate and compensate these delays, since there is usually a possible pulse ordering that does not require strict simultaneity, but only assumes one pulse ends before the other<sup>41</sup>. We enforce this by inserting an additional pulse of 1.4  $\mu$ s duration, which is the minimum pulse length in the control hardware, but long enough to exceed any natural differential delays in the lab.

Another important feature of the DDS hardware is phase tracking [Neg18b; Mar20]. Recall that the Rabi oscillation Hamiltonian in Eq. (2.18) is written in the rotating frame of the drive. Whenever multiple drives require phase-synchronisation, or when a single drive changes frequency over time, the rotating frame needs to be updated. The solution implemented in the DDS hardware is the following. In every shot of the experiment, a certain time is set as  $t = t_0$ . Whenever we wish to apply a pulse with frequency  $\omega$  and phase  $\phi$  at time  $t = T$ , the DDS actually outputs a pulse with phase  $\phi' = \omega(T - t_0) + \phi$ . This way, the desired phase  $\phi$  is applied in the rotating frame of the drive. One consequence of that is that pulses of different frequencies do not have a fixed phase relation. Instead, their relative phase depends on the choice of  $t_0$ ! This is sometimes sub-optimal for sequences of phase-modulated MS gates, as we will discuss further in Sec. 9.2.

### Asynchronous control

For most laboratory equipment, there is no need to synchronise their control with the experimental pulses. Most of these devices are controlled digitally from the control PC. These include:

- PDH laser proportional-integral-derivative (PID) locks (Sec. 4.1.2)
- Fibre phase noise cancellation units (Sec. 4.4.4)

<sup>41</sup> For example, Doppler cooling usually assumes the  $\lambda = 397$  nm and  $\lambda = 866$  nm light is applied simultaneously. If the  $\lambda = 397$  nm is switched off first, the ion will decay into the ground state  $S_{1/2}$ , as it should. However, if the  $\lambda = 866$  nm laser switches off first, the system can become trapped in a meta-stable  $D_{3/2}$  state. For Doppler cooling to function well, we therefore only need to enforce that the repumper switches off last.

### 3. EXPERIMENTAL SETUP

- Trap DACs (recently upgraded to synchronous)
- Slow shutters
- Atomic oven power supply
- Trap heater
- Some laser current/temperature settings.
- Fixed-frequency RF sources

Unfortunately, a number of electrical devices still need to be adjusted manually, although we are in the process of converting them to remote operation:

- Fast lock of  $\lambda = 729$  nm laser (Sec. 4.4.1)
- Injection locks (Sec. 4.4.5)
- Some laser current/temperature settings.

A lot of laboratory devices allow remote control only through serial ports, such as USB. At the start of my PhD, we would daisy-chain those to the control PC through serial adapters. However, this approach was not scalable due to frequent adapter failures and changes in serial port allocation. Other devices only needed a serial port for re-programming, after which they ran stand-alone. This was likewise cumbersome, and I spent hours hunting down hidden devices and their drivers.

Thanks to a lot of work from other group members<sup>42</sup>, we converged onto a better approach. All the remotely controllable devices are now connected to the laboratory LAN and are accessible from any computer in the network. Aside from the control computers, the network comprises a large number of Raspberry Pi units scattered around the lab. Whenever a device requires a USB connection, we plug it into the nearest Pi, upload the drivers, and the device becomes accessible from the network. For most of the communications, we use a custom *plug-in system*, which allows us to quickly set up a webpage for the device interface. The exception is the PID locks, which have a dedicated GUI.

#### Lab monitoring

We use an array of sensors for monitoring laboratory conditions. These include:

- Laser powers
- RF power monitor
- Oven current

---

<sup>42</sup> Most of the credit goes to M. Grau, O. Wipfli and C. Fisher

- Lab temperature, pressure and humidity
- Ion trap temperature
- Dewar helium pressure
- Compressor helium pressure
- Vacuum pressure

All the parameters are logged to an InfluxDB database and displayed in real time on a Grafana dashboard. For devices without a dedicated network or serial connection, we developed a series of circuit boards for plug-and-play monitoring. DC voltages (e.g. output by photodetectors) are connected to small ADC breakout boards<sup>43</sup>, which are powered and interfaced with the lab network through a single ethernet cable connected to an *Ambient box*, developed by O. Wipfli. The Ambient box houses an Arduino microcontroller which logs the digital inputs and sends them to Grafana.

Currently, the weak point of lab monitoring is the laser powers. Due to space constraints, photodetectors could only be added to the existing beam paths in awkward places. That means that it is not trivial to infer the problem when the signal changes. Nonetheless, the monitors do alert you that something has happened and demands your attention.

---

<sup>43</sup> Adafruit ADS1115 breakout board



## 4. Free-space and integrated optics

In this chapter, I present the optical setups used to control and read out the state of the ions. In Sec. 4.1, I introduce the basic tools of the trade. I then present the laser systems used in our lab for loading, cooling, detecting and repumping ions. In Sec. 4.4 I present the  $\lambda = 729 \text{ nm}$  laser used for optical qubit control, and the specific techniques necessary to narrow its linewidth. Sec. 4.5 discusses the free-space light delivery and collection. Finally, Sec. 4.6 discusses the traps with integrated optics and their methods of laser delivery.

### 4.1. Basic laser control

#### 4.1.1. Gaussian beams

The fundamental mode of a laser beam or a rotationally symmetric waveguide is the Gaussian mode [Bro03]. Consider a Gaussian beam of waist radius  $w(z)$  propagating in free space along  $z$ . The transverse intensity profile in the paraxial approximation is given by:

$$I(x, y) \propto e^{-2(x^2+y^2)/w(z)^2}.$$

The beam waist radius  $w(z)$  evolves along the propagation direction  $z$ . Suppose the beam is the smallest at  $z = 0$ , where  $w(0) = w_0$ . The length-scale for beam expansion is called the Rayleigh range and is given by:

$$z_R = \frac{\pi w_0^2}{\lambda},$$

and the waist as a function of propagation direction is given by:

$$w(z) = w_0 \sqrt{1 + (z/z_R)^2}.$$

Throughout this thesis, “beam radius” refers to  $w(z)$  and “beam diameter” means  $2w(z)$ . One implication of the Gaussian beam mathematics is that no beam is truly collimated, and will eventually diverge. To neglect this divergence, we use beams with large  $w_0$  for optical routing. For example, a  $\lambda = 729 \text{ nm}$  beam with  $w_0 = 1.5 \text{ mm}$  has a Rayleigh range of  $z_R = 10 \text{ m}$ .

### Gaussian beam focussing


It is interesting to ask to what extent the familiar ray optics formulae get modified in the Gaussian picture. Consider a Gaussian beam with waist  $w_0$  and Rayleigh range  $z_{R0}$  incident on a thin lens with focal length  $f > 0$ . The beam waist is located at distance  $z_0$  from the lens. Using the Gaussian beam propagation formulae, we can calculate the focus location  $z_1$  of the transformed beam on the other side. The result reads [Sel83]:

$$\frac{z_1}{f} = \frac{\frac{z_0}{f} \left( \frac{z_0}{f} - 1 \right) + \frac{z_{R0}^2}{f^2}}{\left( \frac{z_0}{f} - 1 \right)^2 + \frac{z_{R0}^2}{f^2}}. \quad (4.1)$$

This expression is plotted in Fig. 4.1, together with the geometrical optics prediction:

$$\frac{1}{z_0} + \frac{1}{z_1} = \frac{1}{f}. \quad (4.2)$$

The predictions coincide in two limits. The first is  $z_{R0} \ll f$ , when the input beam waist is



**Figure 4.1:** Focus location  $z_1$  of a Gaussian beam transformed by a lens with a focal length  $f$  versus the focus location  $z_0$  of the input beam. The plot evaluates the expression in Eq. (4.1) for different values of  $z_{R0}$  and compares it to the geometrical optics prediction in Eq. (4.2).

very small and rapidly expanding, and the Eq. (4.1) converges to Eq. (4.2). The second is when  $z_{R0} \gg f$ , and so the input beam is very large and collimated, and we find  $z_1 = f$  regardless of  $z_0$ . This matches the geometrical optics prediction that a collimated beam is focused distance  $f$  away from the lens. In the latter limit (and generally whenever  $z_1 = f$ ), the focused waist size is given by:

$$w_1 = \frac{\lambda f}{\pi w_0}. \quad (4.3)$$



Thus, tightly focused beams can be created by focussing a large input beam through a short focal length lens. Eq. (4.3) specifies a lower bound for achievable spot sizes, and the spots created in practice will be larger unless some care is taken. For a typical singlet lens,  $w_1$  is limited by spherical aberrations, which are more pronounced for shorter focal length lenses. Better performance can be achieved by using multi-element objectives and aspheric lenses. Diffraction-limited aspheric lenses allow us to produce  $w_1$  as low as the diffraction limit, but only in a very small region around the optical axis.

An observant reader will notice a peculiar feature of Eq. 4.1, namely that it is always satisfied by  $z_1 = z_2 = f$ . In other words, a Gaussian beam focused  $z_1 = f$  away from a lens gets re-focused at a distance  $z_2 = f$  on the other side. Doesn't that starkly disagree with the geometrical optics prediction that an image at the focal point of a lens gets imaged onto infinity? Actually, the Gaussian beam calculations clarify the geometrical formula, which exhibits a discontinuous singularity at  $z_1 = f$ . This " $f - f$  configuration" changes a magnification (as given by Eq. (4.3)), unlike the familiar  $2f - 2f$  configuration, which only moves the image without affecting its absolute size [Hec16].

### Small beams in SETs

Suppose we wish to use free-space laser beams to individually address several ions in a one-dimensional crystal in a SET as illustrated in Fig. 4.2 a). In order to minimise cross-talk, we require that the focused beam diameter  $2w_x$  is significantly smaller than the ion-ion spacing  $\Delta x$ . However, rapid expansion of highly focused laser beams poses a challenge to this method, since a beam is likely to clip on the electrodes, causing loss or scatter, as shown Fig. 4.2 b). For a chain of two ions of mass  $m$  in a potential well of axial centre-of-mass frequency  $\omega_{\text{com}}$ , the spacing is given by [Jam97]


$$\Delta x = \left( \frac{2e^2}{4\pi\epsilon_0 m \omega_{\text{com}}^2} \right)^{1/3}$$

where  $e$  is the electron charge and  $\epsilon_0$  is the vacuum permittivity. For a pair of  $^{40}\text{Ca}^+$  ions with  $\omega_{\text{com}} = 2\pi \times 1.2 \text{ MHz}$ , we find  $\Delta x = 5 \mu\text{m}$ , which is a typical ion spacing in this thesis. However, reducing the beam waist to values comparable or below  $\Delta x$  leads to significant clipping even for a short distance  $d$  between the ion and the trap edge, as shown in Fig. 4.2 c).

There are several approaches to this problem. First, tight focusing can be achieved by using elliptical beams, whose intensity profile is given by:

$$I(x, y) \propto e^{-2(x^2/w_x^2 + y^2/w_y^2)}.$$

Ions aligned along the trap axis  $x$  can be then individually addressed by decreasing  $w_x$  below the  $\Delta x$ . On the other hand,  $w_z$  should be made as large as necessary to avoid unwanted scatter. We explored this approach early in my PhD, and a prototype system with



**Figure 4.2:** a) An illustration of the free-space single-ion addressing method in an SET (top view). A focussed pair of laser beams (incident from the left) can be used to address individual ions in a chain (red dots) if  $2w_x \ll \Delta x$ . b) Side view of the challenge associated with the free-space addressing method. A tightly-focussed beam with diameter  $2w_z$  at the ion expands rapidly, leading to scattering on the chip edge distance  $d$  away from the ion. c) Calculated fraction of power clipped on the edge, assuming a beam focused at the ion trapped  $h = 50 \mu\text{m}$  above the surface. The calculation assumes a one-dimensional Gaussian beam.

$w_x = 1.72(3)\mu\text{m}$ ,  $w_y = 6.83(8)\mu\text{m}$  was designed and tested by A. Militaru in his Master's thesis [Mil18a].

The second approach is to not worry about the scatter, and make sure that the light hits the ion before hitting the chip. This can be achieved by addressing the ions through a high-NA objective above the trap. While this method certainly creates danger of significant cross-talk and charging, it might work for certain wavelengths and trap materials. This is also an obvious approach for the SETs with a slot through the electrode surface [Bra16a].

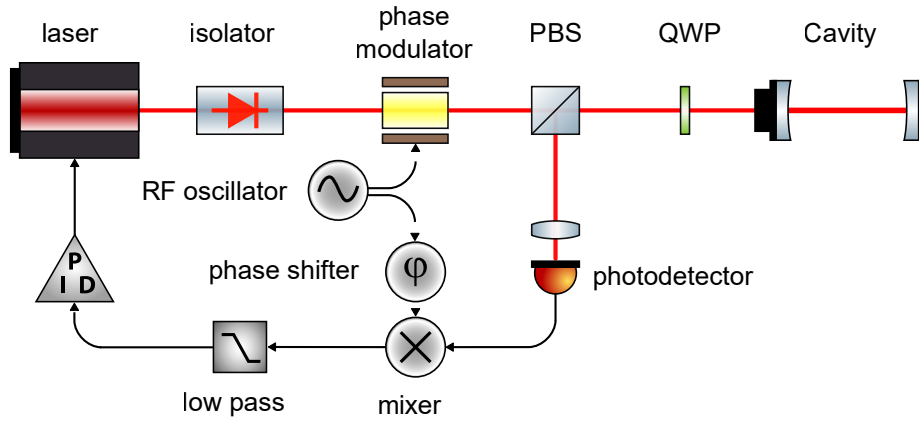
The third approach, which we explored in detail in this thesis, is to send the light from the chip through trap-integrated waveguides [Meh20]. The waveguide is terminated with a diffractive *grating coupler*, which acts as a lens. Since the lens is now very close to the ion, small spots can be created with modest-size diffractive elements, allowing for multi-zone scalability. The clipping is not a worry, since the effective focusing distance is on the order of the ion-electrode spacing.

#### 4.1.2. Frequency stabilisation

An ideal laser would produce light at a single frequency, but in the real world, the frequency is never fixed. We typically distinguish between two timescales of frequency changes.

Frequency changes which are fast compared to the timescale of our operations - typically in the order of  $\mu\text{s}$  - are referred to as the *linewidth*, while slower changes are called *drifts*. A typical laser in our laboratory has a free-running linewidth of  $\Gamma \leq 2\pi \times 500 \text{ kHz}$  at  $5 \mu\text{s}$  time-scale.

As discussed in Sec. 2.3.4, every transition has a certain natural linewidth, and the exact drive linewidth is not important if it is significantly below the natural one. This separates frequency stabilisation requirements into two camps. For dipole transitions, the free-running linewidth is sufficiently low, and we only need to stabilise the laser to avoid frequency drifts. On the other hand, the quadrupole laser at  $\lambda = 729 \text{ nm}$  requires significant linewidth narrowing before high-fidelity qubit operations can be executed.



**Figure 4.3:** General PDH layout. Adapte from an image by Kondephy (CC BY SA 4.0)[Kon16]

Our method of choice for both linewidth narrowing and frequency stabilisation is the PDH lock, where a stable optical cavity is used as a length reference. A detailed description of the technique can be found in multiple references, e.g. [Dre83; Fox03]. Briefly, a portion of the laser beam is picked off and phase-modulated with an electro-optic modulator (EOM). The EOM output is coupled into the reference cavity, and a fast photodetector (PD) records the cavity reflection signal. This signal is then demodulated and passed to a PID controller which feeds back on either the current or the laser grating piezo voltage.

The achievable precision of frequency locking is set by the cavity linewidth and by the speed and resolution of the electronics. For the dipole lasers, the cavity linewidth exceeds the free-running laser linewidth, and slow locking electronics only protect from long-term drifts. For the quadrupole laser, the cavity linewidth is below the free-running linewidth, and fast locking electronics are used to create a much tighter lock. Tab. 4.1 lists the relevant frequency parameters for different lasers in the lab [Lin11]. Further details of the different cavities and frequency locks are discussed in subsequent chapters.


#### 4. FREE-SPACE AND INTEGRATED OPTICS

**Table 4.1:** Laser and cavity linewidths of the sources used in the experiments.

$\lambda(\text{nm})$	Free-running linewidth	PDH cavity linewidth	Locked laser linewidth
397	$< 500 \text{ kHz}$	40 MHz	$< 500 \text{ kHz}$
854	$< 500 \text{ kHz}$	40 MHz	$< 500 \text{ kHz}$
866	$< 500 \text{ kHz}$	40 MHz	$< 500 \text{ kHz}$
729	100 kHz	11.5 kHz	$\approx 10 \text{ Hz} - 100 \text{ Hz}$

##### 4.1.3. Acousto-optic modulation

An AOM is essentially an electrically-activated diffraction grating. An incident RF signal at a frequency  $f_{\text{RF}}$  shakes a piezoelectric transducer, creating a travelling sound-wave with phase velocity  $v_s$  inside the crystal. A laser beam of wavelength  $\lambda$  is incident at an angle  $\theta$  to the acoustic wavevector. When  $\theta = \pm \lambda f_{\text{RF}} / (2v_s)$ , the Bragg condition is satisfied and the light is strongly diffracted (Fig. 4.4). Typically up to  $\approx 85\%$  of power is scattered into the first order, and travels at an angle of  $2\theta = \pm \lambda f_{\text{RF}} / v_s$  with respect to the undiffracted beam [Iso].



**Figure 4.4:** AOM aligned to the +1 order

After a distance  $L$ , where the different diffraction orders are sufficiently separated, an aperture is placed to only transmit the first order beam. This way, the transmitted power can be turned on and off on a time-scale of  $t_{\text{mod}} \approx \phi / v_s$ , where  $\phi$  is the beam diameter. For a typical  $\text{TeO}_2$  crystal (used for red-to-infrared wavelengths)  $v_s = 4260 \text{ m/s}$  resulting in  $t_{\text{mod}} \approx 100 \text{ ns}$  for a  $\phi = 0.5 \text{ mm}$  beam.

In addition to modulating the intensity, an AOM changes the light frequency as well. This is because the diffraction process can be thought of as three-wave mixing between the incident photons and the travelling phonons [Don05]. Following a  $+1/-1$ -order scattering event, the outgoing photon sees its initial energy  $hf_L$  increased/decreased by the energy  $hf_{\text{RF}}$  of the absorbed/emitted phonon. Consequently, the transmitted laser frequency is


given by

$$f'_L = f_L \pm f_{\text{RF}}.$$

where the sign of the shift depends on which diffraction order we consider. We use this process to tune laser frequencies over 10s of MHz.

### Double-pass configurations

Unfortunately, changing the laser frequency by changing  $f_{\text{RF}}$  is not straightforward, since different RF frequencies result in different angles of diffraction, thus modifying subsequent beam-pointing. This severely limits the diffraction bandwidth, since all our lasers are fibre-coupled after frequency modulation. To alleviate this problem, we place the frequency-scanning AOMs in a so-called double-pass configuration, for which one of the setups is shown in Fig. 4.5 [Don05]. When the AOM is turned on, the input beam is diffracted and



**Figure 4.5:** AOM in a double-pass “type A” configuration aligned to +1 order. A quarter-wave plate (QWP) (green) rotates the polarization of the out-going beam with respect to the incident one, allowing them to be separated on a polarising beam splitter (PBS) (left).

incident on a lens which aligns it along the optical axis. The light is then retro-reflected by a mirror, and diffracted once again. By placing the lens  $f$  away from the AOM, we ensure that the beam hits the mirror at the same angle regardless of the angle of diffraction. And by placing the mirror  $f$  away from the lens, we guarantee that an initially collimated beam returns collimated as well. The returning beam is frequency-shifted by  $\pm 2f_{\text{RF}}$  but travels along the same direction regardless of the diffraction frequency.

To separate beams travelling to and out of the double-pass AOM, we employ two different methods. The “type A” configuration, shown in Fig. 4.5, relies on polarization discrimination. The beam travelling towards the AOM is linearly polarized by a PBS. It then passes twice through a QWP aligned to make the returning beam’s polarisation orthogonal to the incident one. As a result, the returning light is separated from the incoming light on the PBS. While this configuration is the most straightforward, it can only be applied when the AOM diffracts both H- and V-polarised light efficiently. This is the case for all of our  $\text{TeO}_2$  AOMs (which pulse the  $\lambda = 729 \text{ nm}$ ,  $\lambda = 854 \text{ nm}$ , and  $\lambda = 866 \text{ nm}$  light). However,

#### 4. FREE-SPACE AND INTEGRATED OPTICS

the fused silica AOMs used for  $\lambda = 397$  nm control only efficiently diffract V-polarized light, and exhibit very large losses otherwise. For those beams, we use what will be referred to as “type B” configuration, where the returning light is still V-polarized but is vertically displaced by a cat-eye retro-reflector and thus returns at a vertical angle to the incident light [Don05; Leu15].

As discussed in Sec. 5.4.1, laser pulses must be well-extinguished to avoid the detrimental effects of light leakage. A double-pass AOM setup can provide sufficient extinction ratios, but problems can be caused by undesired reflections and scattering, as well as beam profile distortions. Ideally, leakage should be measured and corrected if necessary. Appendix 11.4 describes the design principles which minimise the risk of having a problem in the first place.

##### Multi-tone operation

One situation where a double-pass configuration is not desirable is for a multi-frequency operation. Suppose we wish to modulate the beam at two frequencies,  $f_1$  and  $f_2$ , e.g. to drive an MS gate (Sec. 2.4). Upon the first pass, this creates two orders, at frequencies  $f_1/2$  and  $f_2/2$  from the carrier. After the second pass, however, both of the orders undergo another diffraction and travel in the same direction. The outgoing light is now is modulated at three frequencies: in addition to the desired modulation at  $f_1$  and  $f_2$ , half of the power ends up modulated at an undesired frequency  $(f_1 + f_2)/2$ . In addition to losing optical power, this cross-term would create an unwanted carrier drive during an MS gate. Therefore, for multi-tone operation, we only use AOMs in a single-pass configuration. To extend the diffraction bandwidth, we place the AOMs in the Fourier plane between two lenses. The first lens also focuses the incident beam to a small size, reducing the rise time which may be beneficial for fast pulse amplitude modulation. However, focused light is diffracted less efficiently, since a converging beam cannot satisfy the Bragg condition exactly [Iso]. Later in this thesis, we began using single-pass AOMs which are directly fibre-coupled. They are described in Sec. 4.4.7.

##### 4.1.4. Optical fibers

Optical fibres are fundamental to any atomic physics laboratory. In addition to allowing the transmission of laser light between distant locations, they have a practical advantage of splitting the setup into separate “modules”. In this way, optical components can be adjusted before the fibre, and as long as the fibre coupling is re-optimised, the component alignment after the fibre remains unaffected.

### Single-mode fibres

A step-index optical fibre consists of a high-index core surrounded by a low-index cladding. The fibre emits and accepts light at the numerical aperture of:

$$\text{NA} = \sqrt{n_1^2 - n_2^2},$$

where  $n_1$  and  $n_2$  are the indices of refraction of the core and the cladding, respectively. The *normalised frequency parameter*, also known as the V-number, is defined as:

$$V(\lambda) = \frac{2\pi}{\lambda} a \times \text{NA}$$

where  $a$  is the core radius and  $\lambda$  is the guided wavelength [Sny83].

In our setup, we almost exclusively use single-mode (SM) fibres. This increases fibre coupling losses but cleans the beam profile, which means that mode distortions do not accumulate along the beam path. To achieve SM operation, the V-number of a fibre must satisfy:

$$V(\lambda) < 2.405.$$

However,  $V \ll 1$  is also undesirable to limit the optical losses in the cladding. Therefore, each fibre can therefore only support a limited range of wavelengths [Bir97].

### Fibre coupling

Fibre coupling requires the incoming Gaussian beam to match the mode size propagating in the optical fibre. While the NA can be used to calculate the mode size at any wavelength, it is not always a very good approximation. The most precise approach is to deduce the fibre mode diameter  $\text{MDF}(\lambda)$  by measuring the divergence of the out-coupled beam at the desired wavelength  $\lambda$ . A simpler but still pretty good approach is to use a mode diameter measurement at some other wavelength  $\lambda_0$  to extract the effective NA, which is largely wavelength-independent. Some manufacturers specify the effective NA of their fibres directly, while others only provide  $\text{MDF}(\lambda_0)$ , from which the effective NA can be calculated as:

$$\text{NA}_{\text{eff}} = \frac{2\lambda_0}{\pi \times \text{MDF}(\lambda_0)}.$$

The condition for the input gaussian beam of radius  $w_0$  focused by a collimator with focal length  $f$  to match the  $\text{MDF}(\lambda)$  then reads:

$$w_0 = f \times \text{NA}_{\text{eff}},$$

and the same formula applies to fibre outcoupling. Typical SM fibres in our lab have  $\text{NA}_{\text{eff}}$  ranging between 0.07 and 0.10, and we use lenses with  $f \approx 5 - 11$  mm to create collimated

beams with  $w_0 \approx 0.3 - 1.1\text{mm}$ .

### Single-mode Photonic Crystal Fibres

In order to simplify beam delivery, we wish to deliver light at different wavelengths through the same fibre. Light at different wavelengths can be simultaneously guided in a single-mode photonic-crystal fibre (PCF) [Kni96; Bir97]. A PCF cross-section contains a solid-core centre surrounded with air holes running through the cladding. As the wavelength increases and the light becomes less confined, more of the mode overlaps with the hollow cores, which have a lower index of refraction. Consequently, the effective index of refraction decreases with the wavelength, thus keeping the V-number approximately constant. As discussed in Sec. 4.5, this technique allows us to combine all the light necessary for loading and trapping into a single fibre.

### Polarization and PM fibres

An optical fibre distorts the incident polarization due to birefringence. Even in a perfectly symmetric step-index fibre, the stress-induced birefringence causes sufficient phase shifts to effectively randomise the output polarization. To stabilise the polarization, we have two choices: either use a polarisation-maintaining (PM) fibre [Nod86], or carefully stabilise the fibre mechanically [Boc18; Ste19].

A PM fibre is created by deliberately introducing a large degree of birefringence into the fibre core. In this way, the polarization eigenaxes  $x$  and  $y$  are fixed along the propagation direction, and mechanical stress does not couple them together. As a result, light polarized along one of the two eigenaxes remain so, and linear polarization of input light can be faithfully preserved. Define  $P_x$  and  $P_y$  as the fibre output powers after a polarizer aligned along either  $x$  or  $y$ , respectively. The quality of polarization preservation can be then specified through the “polarization extinction ratio”, defined as

$$\text{PER} = \frac{P_x}{P_y}$$

for light aligned to the  $x$ -axis at the fibre input.

We routinely use PM fibres in the lab. Prior to use, one must linearly polarize the input light, and align it to the fibre axis. This can be achieved by minimising the output polarization drifts while mechanically disturbing the fibre<sup>1</sup>. Typically, we achieve polarization extinction ratios of  $\text{PER} \approx 20\text{ dB}$ , although a few fibres in the lab have significantly worse fluctuations.

One empirical observation about PM fibres is that as soon as the beam is not well-aligned to the principal axis, one observes very large polarization fluctuations. I have not systematically gathered any data about this, but my feeling is that a misaligned PM

---

<sup>1</sup> Heating coiled fibre with a heat gun works excellent



fibre introduces more polarization noise than a typical SM fibre. Whether true or not, PM fibres face a challenge when assembled into longer sequences of components, as each subsequent fibre needs to be well-aligned to the previous one. For example, an in-fibre AOM connectorised with PM fibres will generally have a much worse PER than a PM fibre alone, unless the manufacturer does a good job aligning the fibre axes together<sup>2</sup>.

Therefore, in applications where polarization stability is key, the best choice may be to use a standard SM fibre and carefully protect it from external disturbances. The added advantage is that multiple standard fibre components can be then connected. The final polarization can be selected by using a fibre polarization controller, or with waveplates placed before or after the fibre. We use this method to connect the fibre components to our integrated optics trap, which features PM waveguides. This resulted in low noise and excellent long term stability without the need for PM fibres. The only downside is that the components do not naturally return to the same equilibrium, and if the fibres are sufficiently disturbed, the polarization requires realignment.

## 4.2. Photoionization lasers

### 4.2.1. $\lambda = 423$ nm beamline

Originally, the  $\lambda = 423$  nm laser was produced by frequency-doubling a  $\lambda = 846$  nm external-cavity diode laser (ECDL)<sup>3</sup>. The doubling setup<sup>4</sup> was however unreliable. Occasionally, the doubling efficiency would drop significantly, and could only be improved by slightly unscrewing and poking the input fibre. Additionally, the crystal's temperature was sometimes difficult to stabilise, resulting in oscillations of output power.

The latest setup is shown in Fig. 4.6. The  $\lambda = 423$  nm light is produced by a frequency-stabilised diode laser<sup>5</sup>. The light is attenuated<sup>6</sup> to  $\approx 10$  mW and fibre-coupled, and 1% is coupled off with a fibre splitter<sup>7</sup> to the wavemeter for frequency monitoring. The remaining light is then collimated with an adjustable coupler, split between two different experiments, and coupled into a PCF which combines the light with other wavelengths used in the experiment. We typically use  $\approx 160$   $\mu$ W for loading.

### 4.2.2. 375 nm beamline

Originally,  $\lambda = 375$  nm light<sup>8</sup> followed the same beam path as the old  $\lambda = 423$  nm laser. This had several drawbacks. First and foremost, an unacceptable focal shift of the  $\lambda = 375$  nm beam spot at the trap went unnoticed in the original design (Sec. 4.5). We subsequently

<sup>2</sup> Alexander Erhard, private communication

<sup>3</sup> DL 100 from Toptica

<sup>4</sup> Fibre-based unit based on a PPKTP waveguide chip from ADVR


<sup>5</sup> DL pro

<sup>6</sup> At full power, we observed fibre tip degradation

<sup>7</sup> TW470R1A1 from Thorlabs

<sup>8</sup> iBeam smart

#### 4. FREE-SPACE AND INTEGRATED OPTICS



**Figure 4.6:** Photoionization laser beamline. The lasers enter the laboratory and become combined with other dipole beams on the PCF. Shaded:  $\lambda = 375$  nm light temporarily guided to a separate fibre (see main text).

routed laser into a separate fibre, as shown in Fig. 4.6 (semi-transparent). Even after this change, the focal shifts of fibre couplers resulted in low coupling efficiencies and prevented us from increasing loading rates (Sec. 5.1.1)

##### 4.2.3. 388 nm beamline

In a bid to increase the loading rates further, we substituted the  $\lambda = 375$  nm source with a free-running laser diode operating closer to the ionization threshold. This is accomplished with a  $\lambda = 390$  nm laser diode<sup>9</sup>. Since the diode has a  $\pm 2$  nm tolerance, we asked the manufacturer for a diode with a low measured wavelength, noting also that it is often possible to lower the emission wavelength of a laser diode by cooling it [Tob16]. For our particular diode, we measured the wavelength of  $\lambda = 388.32(2)$  nm at a temperature of  $25^\circ\text{C}$  with an optical spectrum analyser (OSA). We also found that the wavelength increases with temperature with a slope of

$$\frac{\Delta\lambda}{\Delta T} = 0.059(2) \text{ nm/K}$$

between  $18.5^\circ\text{C}$  and  $30^\circ\text{C}$ , and that the wavelength is independent of the applied current (above the lasing threshold). We typically operate it at  $20^\circ\text{C}$  and 110 mA. The diode is strongly multi-mode, so we do not monitor it on the wavemeter.

<sup>9</sup> RLT390-50CMG from Roithner Lasertechnik

## 4.3. Dipole lasers

### 4.3.1. Repumper beamlines

The beam paths for repumping lasers at  $\lambda = 866 \text{ nm}$  and  $\lambda = 854 \text{ nm}$  are shown in Fig. 4.7. In some sense, these are the easiest lasers, since their transitions have the lowest saturation intensities (Sec. 2.3.4). The  $\lambda = 866 \text{ nm}$  light is produced by an ECDL<sup>10</sup>, PDH-locked and frequency-shifted by a double-pass AOM<sup>11</sup>. The light is coupled to a PM fibre and delivered onto the optical table in our lab. It is then frequency-shifted once again and split off between different experiments. The beam is pulsed with a double-pass AOM and coupled into a PCF which delivers it to the experiment. The  $\lambda = 854 \text{ nm}$  light follows a very similar pattern: it is produced by a diode laser<sup>12</sup>, coupled into a PM fibre going into our lab, frequency-shifted, split-off between experiments, pulsed with a double-pass AOM, combined with the other repumped with a dichroic beamplitter<sup>13</sup> and coupled into the same PCF. This laser requires extra care to avoid 0th diffraction order light leaking into the fibre (Sec. 5.4.1)

### 4.3.2. 397 nm laser

Light at  $\lambda = 397 \text{ nm}$  is produced by frequency-doubling  $\lambda = 794 \text{ nm}$  light, which is generated by a ECDL and a tapered amplifier (TA)<sup>14</sup>. The  $\lambda = 397 \text{ nm}$  light is coupled into a PM fibre going into our lab. It is then frequency-shifted and split-off between two branches, as shown in Fig. 4.8.

The light on the  $\pi$  branch, used for Doppler cooling and detection, is pulsed with a double-pass AOM, combined with the other repumpers and the photo-ionisation (PI) light<sup>15</sup>, and coupled into the common PCF, which delivers light to the experiment. The light on the  $\sigma$  branch is pulsed with an AOM with a higher central frequency, such that its power is maximal when  $\approx 2\pi \times 60 \text{ MHz}$  detuned from the  $\pi$  light. This creates good conditions for EIT cooling (Sec. 2.5.3.), while still providing some power on resonance for optical pumping. The  $\lambda = 397 \text{ nm}$  laser requires more regular maintenance than any other light source. Typical issues involve:

- Drops of the laser output power, solved by adjusting the alignment of either the  $\lambda = 794 \text{ nm}$  TA seed or the doubling stage. When the position of the doubling crystal is adjusted, the subsequent fibre coupling needs to be optimised as well.

---

<sup>10</sup> DL Pro from Toptica Photonics

<sup>11</sup> This AOM is run with its shutter open, thus producing bichromatic light with a frequency splitting of 240 MHz. Only the (+1,+1) order is useful for us, but the (0,0) order is necessary for another experiment

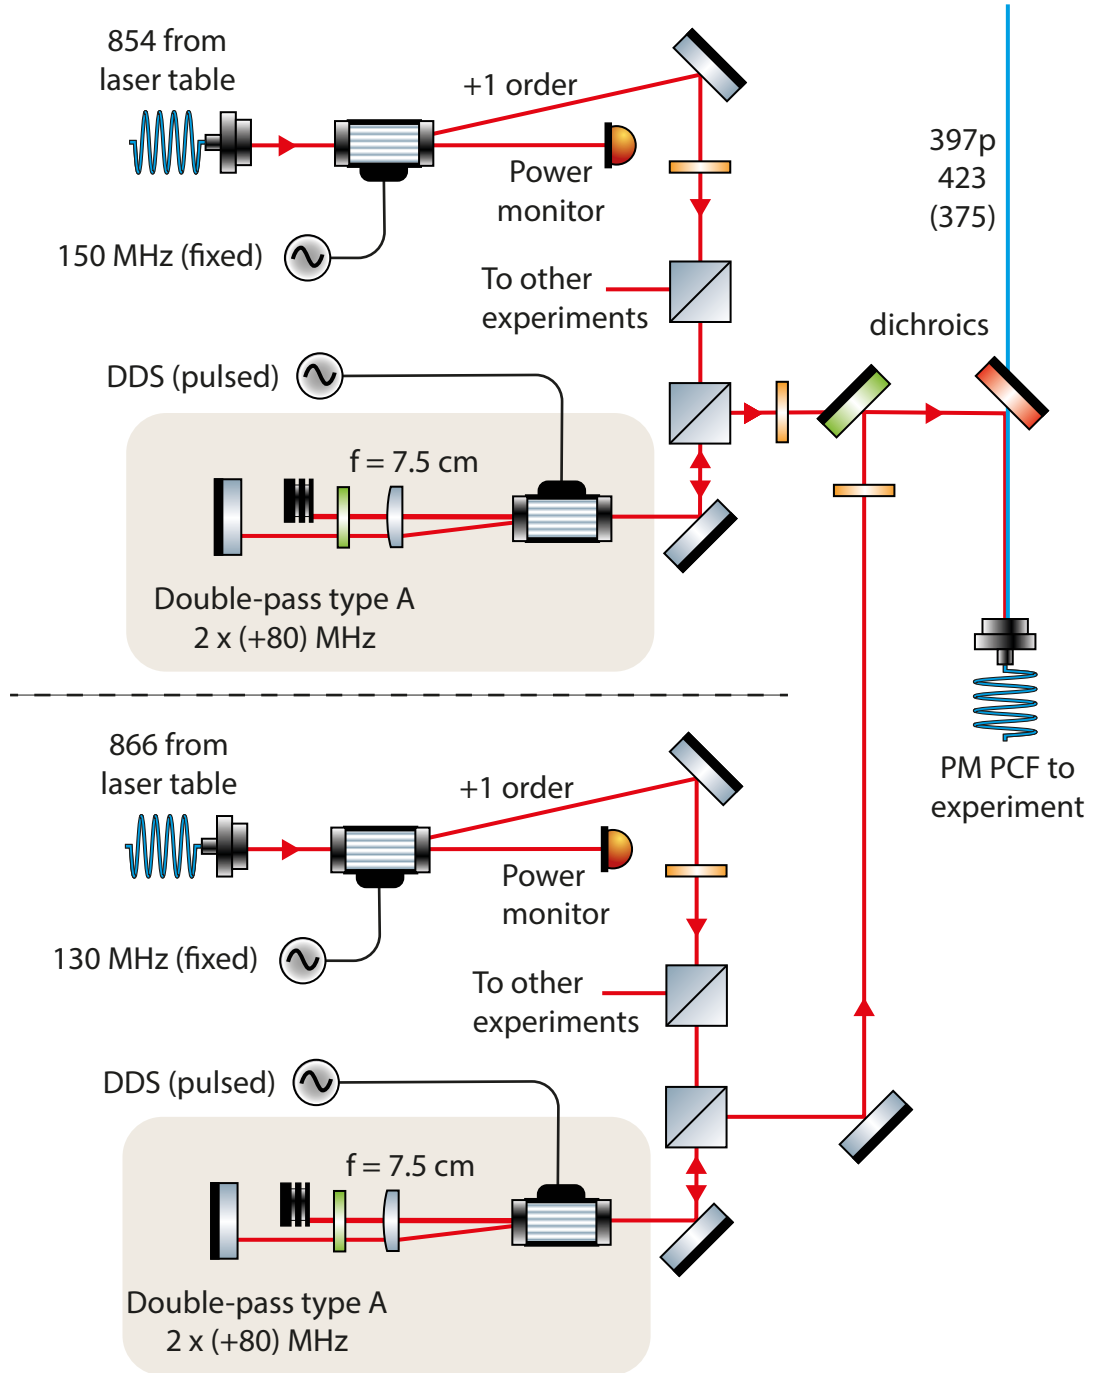
<sup>12</sup> DL Pro from Toptica Photonics

<sup>13</sup> Di02-R442-25x36 from Semrock

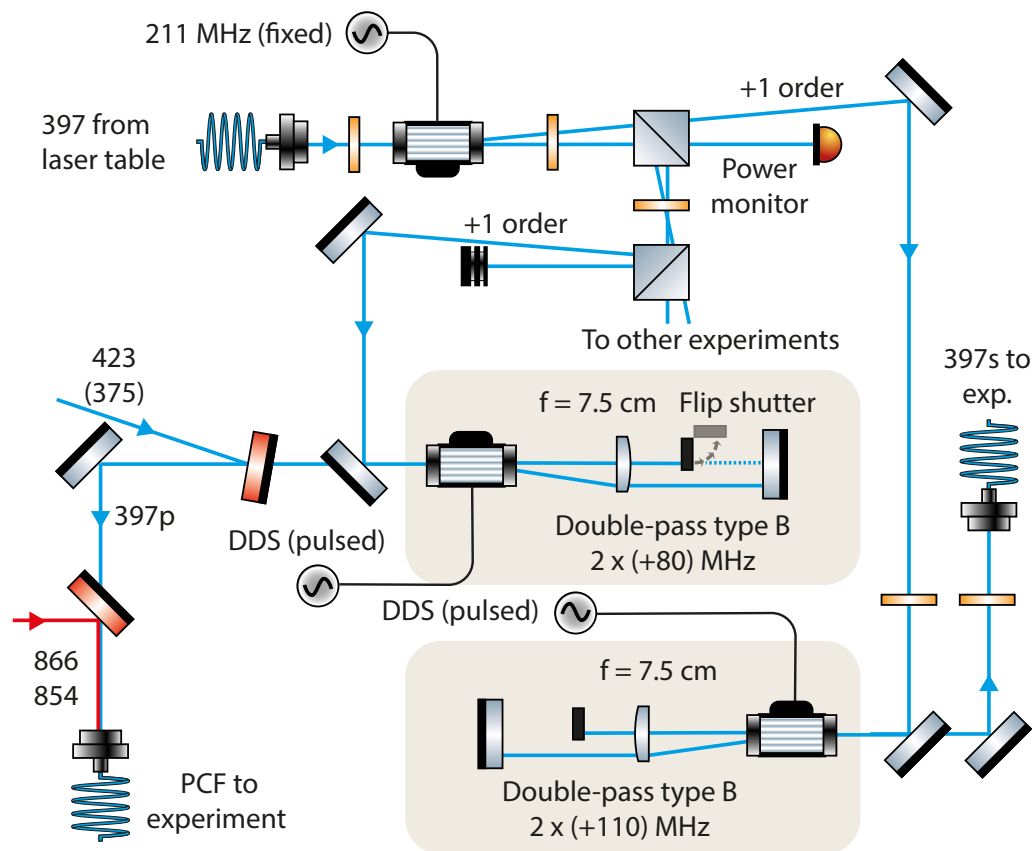
<sup>14</sup> TA-SHG pro from Toptica photonics

<sup>15</sup> Dichroic beamsplitter ZT860lpxr from Chroma Technology is used to combine with the other repumpers, while a bandpass filter FF01-390/18-25 combines the light with the PI lasers

#### 4. FREE-SPACE AND INTEGRATED OPTICS



**Figure 4.7:** Repumper laser setup. Lasers at  $\lambda = 854$  nm and  $\lambda = 866$  nm are individually pulsed and combined on a dichroic, and then combine with other dipole lasers in the PCF.



**Figure 4.8:** 397 laser setup. Light enters through a PM fibre from the laser table, and is split into two branches. One branch provides  $\pi$ -polarized light, which is coupled into a common PCF. The  $\sigma$ -polarized light travels through a separate fibre.

- **Laser-induced damage of the first PBS.** The doubler outputs a 120 mW beam with elliptical Gaussian diameter of 0.68 mm x 0.39 mm. The resulting peak linear power density of  $\approx 6$  W/cm is well below the  $\approx 70$  W/cm damage threshold for the PBS<sup>16</sup>. Nonetheless, the two cubes<sup>17</sup> we used got damaged over time, with optical losses as high as 50% after 1 year of operation. Displacing the cube can eliminate the loss, allowing it to be recycled multiple times. To make the maintenance simpler, we placed this PBS on a translation stage. No damage was observed on PBSs placed further along in the beam path, which handle up to 20 mW of optical power.
- **Polarization fluctuations.** None of the dipole laser beams has their polarisation cleaned

---

<sup>16</sup> PBS25-405-HP high power PBS from Thorlabs

<sup>17</sup> The other one was a custom 397 PBS from Altechna, for which the continuous-wave (CW) damage threshold is not specified

## 4. FREE-SPACE AND INTEGRATED OPTICS

directly after entering our lab via optical fibres<sup>18</sup>. This is a problem for the  $\lambda = 397$  nm light, since the AOM efficiency is polarisation-dependent. Despite introducing an additional waveplate to match the fibre output polarisation to the AOM, the problem remains. The issue is further amplified by the fact that PBS placed right after the AOM is used to split light between different branches. Therefore, as a small drift of fibre polarisation affects both the AOM efficiency and the  $\pi/\sigma$  splitting ratio.

- Laser-induced fibre facet damage, which results in a beam profile distortion. It is solved by re-polishing the fibre face. We only ever observed the damage on the output end of the fibre entering our lab. To maximise the fibre lifetime, we block the light with a TTL-controlled shutter when not in use.

### 4.3.3. Dipole laser frequency control

The dipole lasers are frequency-stabilised using low-finesse cavities, as described in Sec. 4.1.2. The frequency tuning proceeds in three steps. First, the diode current and piezo voltage are tuned until the laser is single-mode and within  $\approx 100$  MHz of the correct frequency, as observed on a wavemeter. Then, each laser is PDH-locked to a separate optical cavity. These in-vacuum temperature-stabilised resonators provide a frequency reference with a typical drift of  $< 5$  MHz over 12 hours. Finally, once an ion is trapped, the laser frequency can be fine-tuned by scanning the cavity length. We adjust the frequency of the  $\lambda = 397$  nm and  $\lambda = 866$  nm lasers by observing ion fluorescence, while the  $\lambda = 854$  nm frequency can be fine-tuned by measuring the repumping speed. While this was never studied systematically, we had some indications that laser frequencies drift more when cavities are scanned more frequently. For this reason, we try to limit cavity adjustments and change laser frequencies with AOMs whenever possible.

## 4.4. Quadrupole laser

### 4.4.1. High-finesse PDH lock

The  $\lambda = 729$  nm laser is generated by an ECDL<sup>19</sup> amplified with a TA<sup>20</sup>. Frequency stabilisation and narrowing are achieved with a PDH lock to a high-finesse cavity. A portion of the diode output is picked off and sent through a phase-stabilised optical fibre (Sec. 4.4.4) to an acoustically-isolated breadboard, which houses the cavity. The reference cavity is a ultra-low expansion (ULE) glass sphere kept close to its “zero-crossing temperature” of  $-4.4^\circ\text{C}$ , where the cavity length is first-order independent of the temperature. Its high finesse of  $F = 270,000$ , combined with a fast locking module with 1 MHz bandwidth<sup>21</sup>, allow

---

<sup>18</sup> This is how the setup was originally designed, and there is no space for modifications

<sup>19</sup> DL pro


<sup>20</sup> BoosTA pro

<sup>21</sup> FALC 110 from Toptica Photonics

us to reach linewidths below 10 Hz [Flü19; Sep12]. The cavity temperature is measured using two 10 k $\Omega$  resistors and controlled with a thermoelectric cooler (TEC) element in a PID loop. The whole breadboard is further cooled to 15 °C to reduce potential cavity temperature gradients. However, the cavity temperature controller has enough cooling power to operate standalone, and the breadboard cooling may even be detrimental due to added vibrations and noise from the chiller<sup>22</sup>. While this warrants further investigations, we have operated with the chiller turned off for a few months, and nobody noticed any difference one way or another. The acoustic noise is suppressed by using a vibration-insensitive cavity design [Lei11] and placing the vacuum chamber in an acoustic enclosure sitting on top of a vibration-isolating platform. Nonetheless, as discussed in Sec. 4.4.8, we have found acoustic noise on the laser frequency, although more experiments are needed to determine whether it originates at the cavity.

#### 4.4.2. Servo bumps

While the PDH lock narrows the linewidth of the main peak, the resulting spectrum is not strictly Lorentzian, as the lock has finite bandwidth. This produces a phenomenon called “servo bumps”, whose influence on the laser spectrum can be measured with an ion, as shown in Fig. 4.9 [Xie17; Fox03]. The servo bump appears around the unity-gain frequency



**Figure 4.9:** Servo bumps seen on the ion in the setup 1.0.

of the controller, where the  $\pi$  phase-shift between the error signal and the correction signal leads to noise amplification, rather than suppression. We can increase the frequency of the servo bump by increasing the lock bandwidth, either by decreasing the hardware response time (through shorter cables, faster PDs etc.) or by increasing the lock gain. Nonetheless, servo bumps pose a serious limitation to coherent sideband operations, since it is difficult to ensure they are decoupled from the motional sidebands, especially for multi-ion crystals.

<sup>22</sup> Not to mention these chillers tend to leak and are loud

While the light with servo bumps was used for all the experiments in setup 1.0, we later upgraded the setup to obtain a cleaner spectrum for multi-qubit quantum logic.

### 4.4.3. Upgraded quadrupole beamline

To improve spectral purity, a new path was created, with the high-finesse cavity used as a spectral filter, extending on the work described in [Flü18; Ger15]. This beamline is now used for coherent waveguided operations, while the original light (referred to as the “debug 729”) is occasionally used for probing micromotion.


The new beamline is shown in Fig. 4.10. A small amount ( $\approx 10 \mu\text{W}$ ) of spectrally filtered light is transmitted through the high-finesse cavity. The light is used to injection-lock a diode, which amplifies the power to  $\approx 40 \text{ mW}$ , serving two experiments. After travelling through a phase-stabilised fibre,  $\approx 2 \text{ mW}$  are used to injection-lock another diode, while the rest is sent to other experiments. The injected output seeds a TA, whose  $\approx 400 \text{ mW}$  output is fibre-coupled and sent to the AOM stage. Pulsing is achieved with a double-pass AOM, while an in-fibre single-pass AOM is used to provide multiple tones for MS gates. In the subsequent sections, we discuss in detail different sections of the upgraded beamline.

### 4.4.4. Fibre noise cancellation


Long optical fibres introduce acoustic phase noise onto the light, resulting in increased linewidth. To perform high-quality coherent operations, the laser phase needs to be stabilised. The components we use for that purpose have been developed in-house and are described in [Opp15]. Here I give a brief overview of the principle, discuss the operational challenges and their solutions.

The fibre noise cancellation (FNC) system is based on a beat-note measurement between a signal beam, which measures the fibre phase distortion, and a reference beam [Ma94]. Our implementation of an FNC system is shown in Fig. 4.11. The signal beam is diffracted in  $+1$  order on an AOM before entering a fibre. We aim to stabilise the phase  $\phi$  of light leaving the fibre, which is a sum of phase  $\phi_{\text{AOM}}$  imprinted by the AOM and  $\phi_{\text{fibre}}$  introduced by propagation in the fibre. A small fraction of light is retro-reflected and travels back through the fibre, accumulating extra phase  $\phi'_{\text{fibre}}$ . For acoustic and mechanical distortions, the time-scale of variations (up to  $\approx 50 \mu\text{s}$ ) is significantly slower than light round-trip time ( $100 \text{ ns}$  for a  $15 \text{ m}$  fibre), therefore it is safe to assume  $\phi'_{\text{fibre}} = \phi_{\text{fibre}}$ . The returning beam is then passed again through the AOM, and the undiffracted order is aligned on a PD. The reference beam is transmitted through the AOM, retro-reflected, and diffracted on the AOM in the  $-1$  order. As a result, the PD records a heterodyne signal with frequency difference  $2f_{\text{AOM}}$  and phase difference  $2\phi = 2(\phi_{\text{AOM}} + \phi_{\text{fibre}})$ . The PD output is AC-coupled to remove the DC component, amplified and mixed with a local oscillator at  $2f_{\text{AOM}}$  to extract the phase difference  $2\phi$ . The phase information is then fed to a PID controller, which adjusts a voltage-controlled oscillator (VCO) to keep  $\phi$  at a desired setpoint.





**Figure 4.10:** Narrow-linewidth  $\lambda = 729$  nm laser beam path



**Figure 4.11:** A FNC system. Light travelling through the fibre and back is combined with a reference signal. The resulting heterodyne beat-note contains the phase information about the fibre phase distortion, which is fed back onto the AOM. POISON is the board used for the fibre-noise cancellation, EVIL is the lab PID control unit, and LO stands for the local oscillator.

### Photodetectors

A high-bandwidth detector is necessary to detect oscillations at  $2f_{\text{AOM}}$ . Our PDs are based around reverse-biased S5973-series photodiodes from Hamamatsu, whose bandwidth exceeds 1 GHz. Our homemade build is similar to the KPD110 circuit described in [Kei07]. Based on the measurements of the original circuit, we expected the gain of  $\approx 200 \text{ V/W}$ . Unfortunately, different detectors used in the lab were built by different people at different times, usually without adequate documentation. Consequently, the level of noise and the gain vary significantly between supposedly identical boards. One cause of this - identified by R. Oswald - was a mistake in voltage regulator filtering<sup>23</sup>, and fixing it removed a lot of noise spikes. Still, lacking a reproducible and large signal, we started replacing the PDs with ET-2030A boards from EOT. These budget detectors produced a significantly larger signal (by  $\approx 10 \text{ dB}$ ), while keeping the noise comparable, thus improving the signal quality overall. Given their specified gain of  $450 \text{ V/W}$  and noise-equivalent power (NEP) of  $< 60 \text{ pW}/\sqrt{\text{Hz}}$ , we can approximately spec our home-built KPD PDs at a gain of  $\approx 120 \text{ V/W}$  and NEP  $< 60 \text{ pW}/\sqrt{\text{Hz}}$  in the frequency range of  $200 \text{ MHz} - 350 \text{ MHz}$ . All FNCs in our lab are now being upgraded with ET-2030A detectors, while KPDs are still used for PDH locks.

### Phase measurement and locking

While some of the FNC locks are almost unconditionally stable, others come out of lock frequently. Another odd feature is that some locks need to be locked to a setpoint of  $\approx 0$ , or they become unreliable otherwise. While writing this thesis, I discovered a likely explanation of both issues. In order to measure the RF phase of the PD signal, we use

<sup>23</sup> It seems likely that some people in the group realised it before. I once removed and opened one particularly well-performing detector, and found that its voltage regulator was removed and bypassed with a wire.

a mixer as a de-modulator. Suppose a signal  $E_1 = \Re(|A_1|e^{i\omega t + \phi})$  is mixed with a local oscillator  $E_2 = \Re(|A_2|e^{i\omega t})$ . The output signal is then filtered by a low-pass filter with cut-off frequency  $\ll \omega$ . The output is then

$$E_3 = k|A_1||A_2|\cos(\phi).$$

Measuring  $E_3$  allows us to extract  $\phi$ , but there is a danger in this mode of operation since the value of the signal depends on  $|A_1|$  and  $|A_2|$ . While we can usually ensure that the strength  $|A_2|$  of the local oscillator is constant, we often don't have the same certainty about  $|A_1|$ . Fast changes of  $\phi$  may change the AOM diffraction angle sufficiently to affect  $|A_1|$ , causing a feedback loop failure. This explains why the only stable point on some locks is  $E_3 = 0$ , as that is the only point where variations of  $|A_1|$  do not affect the error signal,

The simplest solution is to run the mixer close to saturation [Rub06]. Because of mixer non-idealities, the output power saturates when both  $|A_1|$  and  $|A_2|$  are of sufficient magnitude, and then the output depends on  $\cos(\phi)$  only. We are now in the process of amplifying our PDs to ensure mixer saturation.

### Back-reflection and isolation

In our FNC systems, back-reflections are generated in one of two ways. In the first method, we use flat FC/PC connector on the fibre output. The fraction  $r$  of back-reflected power can be estimated from Fresnel equations for a flat dielectric interface:

$$r = \left( \frac{n_{\text{fibre}} - n_{\text{air}}}{n_{\text{fibre}} + n_{\text{air}}} \right)^2.$$

For a fused silica core at  $\lambda = 729 \text{ nm}$ ,  $n_{\text{fibre}} = 1.45$  and  $n_{\text{air}} = 1.0$ , leading to  $r = 3\%$ <sup>24</sup>. This method of generating back-reflections is the simplest and usually provides a sufficient signal. When additional power is required, we use a second method, whereby a partially reflective slide is inserted after the fibre collimator. By choosing the right retro-reflector, we can achieve any desired ratio between the transmitted and the reflected power<sup>25</sup>. Particular care must be taken to ensure that all powerful reflections at the correct frequency only come from the retro-reflector, and not from any subsequent elements. I give some tips about dealing with back-reflections in Appendix 11.3.


#### 4.4.5. Diode injection locking

An important method for amplifying CW lasers is by injection locking, whereby a laser diode is forced to emit at the frequency set by a weaker seeding laser [Had86]. A diode injection system is schematically depicted in Fig. 4.12. The usual nomenclature for the

<sup>24</sup> This estimate is consistent with the  $-14 \text{ dB}$  return loss specified for their flat end-face connectors by OZ optics.

<sup>25</sup> But beware! This thing is tricky to align

#### 4. FREE-SPACE AND INTEGRATED OPTICS



**Figure 4.12:** Schematic of an injection lock. Parent light seeds a child diode, which is current-controlled and held in a temperature-controlled mount. At specific values of the diode current, the child output becomes locked to the parent frequency. The parent and child beams are separated with an optical isolator.

parties involved is *Master* and *Slave*. I think a better way of describing this process - less colonial and still conceptually illuminating - is to call them *Parent* and *Child*. The *Parent* is a spectrally pure laser whose optical power is not sufficient for the experiment. The *Child* diode is sufficiently powerful, but, when free-running, does not have the necessary spectral purity. By sending the *Parent* light at the *Child* diode, and overlapping their modes, the *Child* diode can become injected, thus emitting most of its power within the lineshape of the *Parent*. Practical instructions on setting up and optimising injection locks are given in Appendix. 11.2.

Our setup comprises two injection locks in sequence. The first diode<sup>26</sup> is placed in the acoustically isolated cavity box, housed in a homemade mount and controlled by a low-noise laser controller<sup>27</sup>. The diode is maintained at  $T = 35.8^\circ\text{C}$ , which was found to be the optimal temperature to amplify a small amount of seed power ( $\approx 10\ \mu\text{W}$ ) [Ger15]. The second diode<sup>28</sup> is in a more noisy environment, but it is operated with  $\approx 2\ \text{mW}$  seed power. It is housed on a temperature-controlled mount<sup>29</sup>, driven by a Thorlabs current controller<sup>30</sup> and a Toptica temperature controller<sup>31</sup>. Both injections use the same large-aperture Faraday isolator<sup>32</sup> and the diode collimation is achieved with a symmetric aspherical lens in an adjustable mount. A portion of the child output is picked off and sent to a scanning Fabry-Perot interferometer for continuous linewidth monitoring.

<sup>26</sup> HL7302MG from Thorlabs

<sup>27</sup> Vescent photonics D2-105

<sup>28</sup> Initially AR-coated LD-0730-0040-AR-2 diode from Toptica Photonics. Later, HL7302MG from Thorlabs

<sup>29</sup> LDM56 from Thorlabs

<sup>30</sup> LDC202C

<sup>31</sup> DTC-110

<sup>32</sup> Qioptiq LINOS FI-730-5SV

### Injection quality

When setting up and optimising injection locks, I assessed their quality using several qualitative and quantitative measures. An ideal injection lock:

1. Produces child spectrum indistinguishable from the spectrum of the parent laser.
2. Can be operated with low seed power.
3. Stays injected indefinitely, without the need for adjustments.
4. Stays injected regardless of drifts in seed power, current or temperature.
5. Produces sufficient output power.

Generally speaking, injection locks are more stable at low currents, and less stable at high currents (where the output power is highest as well). The factors that affect high-power injection stability are:

1. **Diode coating.** We have extensively experimented with two versions of the same laser diode, one with<sup>33</sup> and one without<sup>34</sup> anti-reflection (AR) coating. The coated diodes were consistently more stable, and the injected spectrum matched the Parent spectrum over a larger range of currents and temperatures [San18]. This is expected since AR coating increases the laser cavity linewidth [Had86].
2. **Diode gain profile.** Just like children, no two diodes are created equal, and we attribute differences between nominally identical diodes to the differences in their gain profiles. The variability can be significant and persists over a range of temperatures. For example, for three different AR-coated diodes in the same mount, injected at the same free-running power with the same seed power and the same mode-matching method, I observed  $\approx 40\%$  variability in the injected output power.
3. **Seed input power.** Reducing the seed input power reduces the range of currents for which the diodes can be injected. For example, our AR-coated  $\lambda = 729$  nm diode was unconditionally injected at every current  $< 100$  mA when seeded with 3 mW of light. Reducing the seed power 10-fold to 0.3 mW decreased the clean injection range to  $\pm 1$  mA around the optimal points [San18].
4. **Mode-matching.** Since laser diodes produce elliptically-shaped beams, additional mode-shaping optics are necessary for mode-matching with a single-mode Gaussian beam. I have never explored this route, and instead simply tried to ensure that the input beam diameter is close to the average diode output diameter. While such mode-matching generally improved the injection quality, the gains have been modest.

---

<sup>33</sup> LD-0730-0040-AR-2 from Toptica Photonics

<sup>34</sup> HL7302MG from Thorlabs


#### 4. FREE-SPACE AND INTEGRATED OPTICS

A much more relevant factor was the size of the input beam. I have found a sharp drop in the injection quality when the seed beam was decreased in size from  $\phi \approx 1.5$  mm to  $\phi \approx 1.0$  mm, regardless of the diode output beam size. I speculate that the mode-matching error is dominated by diode output mode distortions. A small seed beam requires shorter diode collimation optics, which introduce more aberrations, further accentuating the beam profile error.

5. **Back-reflections.** While the isolator used for injection locking provides a large degree of isolation from further back-reflection, it can be insufficient when a strong back-reflection is present, causing instability.

#### Injection noise

Initially, found that both injection locks produced significant non-white amplitude noise. On the first diode, I found a ripple, repeating every  $33 \mu\text{s}$ , with an amplitude decreasing for higher currents (Fig 4.13, left). This ripple was eliminated in two steps (Fig 4.13, right).




**Figure 4.13:** (left) Measured ripple amplitude vs diode current at the injection points. The results are connected by a line to guide the eye. (right) The reduction in the measured injected ripple size at 81 mA. The final dippler is beyond the shot-noise background.

First, its amplitude decreased after swapping a Toptica current controller<sup>35</sup> for a Vescent controller, presumably because the latter supplies a differential output, and thus does not introduce a ground loop when combined with a temperature controller. Secondly, the ripple was reduced below the noise measurement background by connecting the diode housing to the breadboard earth through a capacitor.

On the second injection setup, the dominant ripple repeated every 50 ms. It turned out that the current controller would spike whenever the temperature controller in the same rack was connected to the diode mount. Replacing the current controller with a stand-alone unit removed any observable ripples when injected (Fig 4.14, right). I further optimised the electrical noise by monitoring the optical output of a free-running diode (Fig 4.14, left).

<sup>35</sup> DCC-110



**Figure 4.14:** (left) Measured ripple amplitude for a free-running diode at 75 mA. At the end, some ripple is visible over the measurement background, but it cannot be seen on the injected output. (right) The reduction in the measured injected ripple size at 75 mA. The final ripple is beyond the shot-noise background.

I tried to float the diode housing and supply the current differentially, but the increased pick-up on the unshielded homemade cable adapter only increased the problem. Finally, explicitly grounding the diode mount to the breadboard earth turned out to be the best solution, though some noise remained.

#### Death of diodes

Originally, an AR-coated diode was used in the second injection stage, which produced a very stable lock. Unfortunately, it began to deteriorate after about 12 months of use. A replacement AR-coated diode was installed, but this one did not last long either and began to deteriorate after 6 months already. In both cases, this would manifest itself by a sudden drop of the output power (by about 30-50%), followed a slower decrease in output power over weeks. The third AR-coated diode we had available was unsuitable, as it outputted 40% less power than the previous two when injected. It is not clear whether the deaths were caused by misuse or not. While the diodes were never operated above the maximum allowed current, both were at some point operated without proper cooling<sup>a</sup>.

At that point, we swapped the diode to a non-coated one. This means the lock is significantly less stable, and the current requires re-adjustment approximately once per day if the lab is busy (and very rarely if there is nobody around). On the other hand, the diode has not deteriorated in performance over the last two years. The non-coated diode was found to be more easily disturbed by back-reflections. Despite the protection of a Faraday isolator, we found that back-reflections from the TA disturb the child spectrum. A second isolator<sup>b</sup> was therefore installed between the injection and the amplifier. This raised the total isolation from 40 dB to 80 dB, and no further disturbance could be observed.

<sup>a</sup> In the case of the first diode, the temperature controller broke without anyone noticing. The injection was so stable that we nonetheless operated like this for months. The second diode was used for about a day without proper heat-sinking due to an assembly error.

<sup>b</sup> IO-5-730-HP from Thorlabs

#### 4.4.6. Tapered amplifier

The output power of a diode injection lock is limited by the maximum power handling of the diode. At high optical powers, diode failure can occur due to facet damage or electrical degradation from high currents. Our  $\lambda = 729$  nm diodes allow us to produce up to 40 mW of single-mode light. Further gain can be achieved by using a TA. A TA chip is essentially a laser diode, AR-coated on both ends, combined with a waveguide. The waveguide is tapered to allow the beam to gradually expand as its power grows. This way, high amplification can be achieved while preserving beam quality and avoiding damage. The TA is injection-locked in a single-pass configuration: the parent light comes in on one end, and leaves from the other end. Our TA is a commercial unit from Toptica<sup>36</sup>, with a pre-installed output isolator and a cylindrical output fibre collimation lenses. Initially, the output power (seeded or not)




**Figure 4.15:** Photo of a typical TA output ripple as seen on the oscilloscope (green trace). The TA output is attenuated and sent of a PD, which is AC-coupled on the oscilloscope.

exhibited a pronounced ripple (Fig. 4.15). Close to the maximum power, the seeded output power fluctuations were at the level of 2%. Eventually, I connected the TA to a different current controller<sup>37</sup>, and the problem went away (Fig. 4.16, left). The seeded output power depends both on the current and on the power in the seed beam (Fig. 4.16 right). Initially,

<sup>36</sup> boostTA pro

<sup>37</sup> DCC-110 1000 mA current controller





**Figure 4.16:** Seeded TA performance. (left) Ripple voltage vs current. Ripple disappears by replacing the manufacturer-provided current controller with a generic one. (right) The TA output power for different values of seed power. The maximum achievable power nowadays is  $\approx 400$  mW is less than the original output of  $\approx 500$  mW, likely just due to a decrease of available seed power.

we were able to obtain over 550 mW of output power with 40 mW seed power. Since the latest diode requires additional isolation from the TA reflections, less seed power is available, and we are typically limited to 400 mW output power.

The TA output is directly coupled to a PM fibre by a manufacturer-supplied unit<sup>38</sup>. This coupler is limited to  $\approx 50\%$  efficiency and does not fully secure the fibre. Consequently, after applying a significant mechanical force on the fibre, the coupling has to be re-optimized. The fibre output is passed through a PBS and reflected by a backside-polished mirror with  $\approx 1\%$  of the power transmitted and fed onto a PD. The PD output is passed to a PID controller which feeds back on the TA current to stabilise the output power with up to 1 MHz bandwidth.

#### 4.4.7. Pulsing and multi-tone generation

The CW output of the TA is fed to a pulsing stage with a double-pass AOM<sup>39</sup> and an in-fibre single-pass AOM<sup>40</sup>, which is then directly connected to the trap-integrated waveguides (Sec. 4.6.4). The in-fibre AOM performs very well. The total device efficiency is  $\approx 42\%$ , which is a product of our fibre coupling efficiency ( $\approx 60\%$ ) and the intrinsic unit efficiency ( $\approx 70\%$ ). At the centre frequency of 147 MHz, we measure a diffraction bandwidth of 16 MHz (FWHM), high extinction ratio of 55 dB<sup>41</sup>, rise and fall time of 30-35 ns, and negligible polarization dependence.

<sup>38</sup> Fiberdock from Toptica Photonics


<sup>39</sup> ATM-1601A2 from Intraaction

<sup>40</sup> Custom unit for  $\lambda = 729$  nm, based on SFO4733-T-M150-0.5C2W-3-F2S FibreQ-AOM from Gooch & Housego

<sup>41</sup> At the time of this measurement, we recorded an unknown noise at 150 MHz being picked up by all the cables in our laboratory. This reduced the extinction ratio to 47 dB when the AOM is connected to the RF amplifier.

#### 4. FREE-SPACE AND INTEGRATED OPTICS

The double-pass AOM is used to scan the frequency, with the single-pass typically operated at a fixed centre frequency, since it has a more limited diffraction bandwidth. We apply single-tone laser pulses by first switching on the single-pass to the desired amplitude, and then using the double-pass to smoothly ramp the pulse on- and off.



**Figure 4.17:** AOM nonlinearity measurement. The measurement is performed by feeding tones at  $f_0 \pm f$  with equal amplitudes to the single-pass AOM. The component at  $2f$  is the desired response, and higher even-order components indicate a presence of undesired frequencies  $f_0 \pm 3f$ .

Multi-tone operation is implemented by adding two DDS tones which are symmetrically detuned from the centre frequency. Both DDS tones are simultaneously fed to the single-pass AOM, while the double-pass sets the centre frequency and the pulse shape. We then send the output to the photodiode and adjust the power in both sidebands. As the RF power in the two tones increases, we begin to observe significant AOM nonlinearities. At a certain point, increasing the total RF power increases the overall diffraction efficiency, but the power at the desired tones drops. This can be measured by feeding the AOM with two RF tones at  $\pm f$  from the centre frequency, and monitoring the strength of the beat-note at the frequency of  $2f$ , as shown in Fig. 4.17. The useful power can be modestly increased by feeding higher-order tones to the AOM at specific phases [Sta20], but we haven't yet applied this technique to coherent operations. Instead, we operate the AOM at the point where the power in the desired orders is maximal.


##### 4.4.8. Laser spectrum

A common method of laser spectrum measurement is with a self-homodyne setup, where a beat-note between two paths with different delays reveals the total linewidth [Lud98]. In such a measurement, the differential delay length must exceed the coherence length of the laser. This makes the method impractical for our  $\lambda = 729$  nm light, whose expected linewidth  $\Delta f \approx 10$  Hz requires delays in excess of  $L = c/(\Delta f) \approx 30\,000$  km.

In order to get information about the linewidth, we instead beat together two separate narrow-linewidth  $\lambda = 729$  nm lasers from different laboratories. We can then extract the approximate linewidth by assuming both sources have the same spectrum. We performed such measurements twice, although never on the cavity transmission. This informs us about the linewidth we can expect from the PDH lock and the FNC systems alone, but not about the injection quality.

The measurements performed so far are inconclusive. The first measurement, performed by C. Fluhmann and described in her thesis [Flü19], recorded a beat-note linewidth  $\Delta\omega < 2\pi \times 10$  Hz FWHM for an averaging time  $t < 100$  ms, limited by the spectrum analyser resolution. She also recorded additional noise peaks at  $2\pi \times 50$  Hz and  $2\pi \times 180$  Hz. Afterwards, the ancillary laser system in the other laboratory was moved and re-built. Recently, the measurements were repeated, and we recorded a beat-note linewidth of  $\Delta\omega = 2\pi \times 50$  Hz FWHM for the same averaging time, and a strong noise peak around  $2\pi \times 185$  Hz. Further investigation is needed to establish which of the two lasers increased in linewidth, but we have evidence that the noise peak is only present on our laser.

Information about the linewidth can be obtained by Ramsey measurements (Sec. 5.4.2). As will be discussed in Sec. 5.4.3, we observe a frequency fluctuations at the level of  $\Delta\Omega \approx 2\pi \times 20$  Hz superimposed with strong discrete noise at a frequency of  $\approx 2\pi \times 175$  Hz and amplitude of  $\approx 2\pi \times 185$  Hz. Further experiments are necessary to find out whether or not injection locking increases the laser linewidth.



**Figure 4.18:** Off-resonant scattering rate vs detuning from the nearest carrier transitions. Each dataset is fitted with a linear model with a constant excitation rate (solid), and the fitted excitation rate is displayed in the legend. e For the near-resonant datasets (red, blue, yellow), we ensure that only one carrier transition is close to resonance and that the carrier Rabi frequency exceeds the sideband Rabi frequency. The curves are normalised to the carrier Rabi frequency of  $\Omega = 2\pi \times 140$  kHz. The far-detuned dataset (green) is obtained by locking the cavity one free spectral range away. The ion is driven with the same optical power, and the curve is unnormalised.

One specific concern is whether the spontaneous emission of the laser diode is effectively

suppressed by the injection lock, or whether the noise floor becomes amplified upon injection. We began investigating this effect by measuring the probability of  $\lambda = 729$  nm exciting the ion when it's far off-resonance from any carrier transition. Fig. 4.18 summarises these preliminary observations. The laser power corresponds to a carrier Rabi frequency  $\Omega = 2\pi \times 140$  kHz on the  $|-1/2\rangle \rightarrow |-1/2\rangle$  optical transition. At a detuning of  $\approx 2\pi \times 2.2$  MHz from this carrier, the off-resonant drive excites the optical qubit at a rate fitted to be  $\approx 10$  Hz. I speculate that this is dominated by the excitation of the  $|-1/2\rangle \rightarrow |-1/2\rangle$  carrier, since all sideband transitions are detuned by  $> 2\pi \times 1$  MHz, and other carriers by  $> 2\pi \times 7.9$  MHz. I repeated the measurements for larger detunings, always ensuring one carrier transition is significantly closer than all others. Fig. 4.18 shows the measured scattering rates (normalised by the carrier Rabi frequency). We find that the scattering rate drops by a factor of 4 between  $\approx 2\pi \times 2$  MHz to  $\approx 2\pi \times 20$  MHz. Finally, by locking the laser cavity to a mode detuned by  $2\pi \times 3.5$  GHz, we measure that the scattering rate is reduced by a further factor of 3.

This measurement is just the first step towards understanding the injected spectrum. Future experiments should compare the scattering rate between the upgraded beamline (with two injection locks) and the original beamline (with no injection locks). This will allow for qualitative evaluation of the spectral impurity. In addition, further investigations are needed to evaluate the coherence of the spectrum.

Finally, we have observed significant spectral impurity from badly injected diodes. More worryingly, there are times when the diode output spectrum appears clean on the monitoring Fabry-Perot interferometer<sup>42</sup>, but the rate of off-resonant qubit excitation is vastly increased. In future experiments, it is worth developing a more sensitive monitor of the injection quality, whether through better optical measurements or using the ion as a probe.

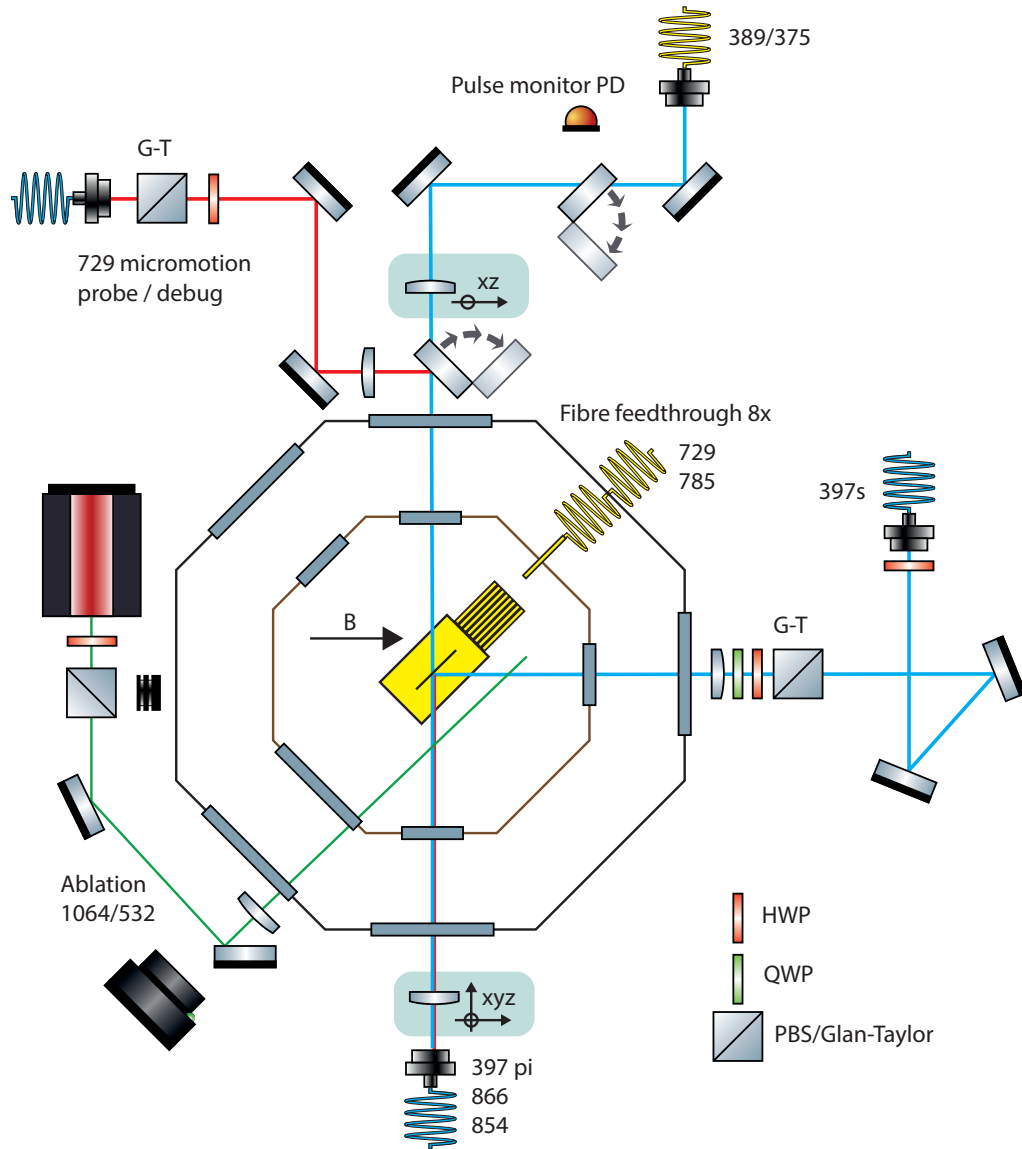
## 4.5. Light delivery and collection

### 4.5.1. Free-space light delivery

Fig 4.19 presents an overview of beam delivery to the ion trap. Lasers at 397 ( $\pi$ -polarized), 423, 854 and 866 nm are all guided through a single PM PCF fibre oriented at 90 deg to the quantization axis.  $\lambda = 397$  nm  $\sigma$ -polarized light is guided through a separate fibre which aligns it with the B-field, and the PI light at  $\lambda = 375$  nm or  $\lambda = 388$  nm was moved from the PCF to a separate fibre to avoid focal shifts. A free-space  $\lambda = 729$  nm beam, occasionally used for debugging, is delivered through the same port as the PI light. An ablation laser travels through a separate viewport and hits the target located next to the trap chip.

---

<sup>42</sup> Scanning Fabry-Perot interferometer SA200-5B from Thorlabs, confusingly referred to in other theses from our group as an “OSA”




**Figure 4.19:** Beam delivery layout on the optical table (viewed from above).

## PCF

The core element of experimental light delivery is a polarization-maintaining PCF<sup>43</sup>. It combines all the wavelengths necessary for trapping (and initially also for loading) into a single package that just requires one-time alignment. The light from the fibre is out-coupled by a chromatically-corrected microscope objective<sup>44</sup> and focused onto the ion with a  $f = 200$  mm achromatic doublet lens<sup>45</sup>. The assembly is mounted on a manual three-axis translation stage<sup>46</sup>, and beam tilt can be controlled by displacing the focussing lens transversely to the optical axis<sup>47</sup>. All the wavelengths are focused to below  $30\text{ }\mu\text{m}$  waist radius.

For a long time, we believed that the total focal shift across all wavelengths was 5 mm. However, after many unsuccessful loading attempts in the integrated waveguide traps, we found that a dramatic focal shift of the 375 nm laser went unnoticed for years (Fig. 4.20). These measurements revealed that we were likely always scattering significant power of UV light off trap electrodes. The focal shift was found to originate already at the collimator, and so we subsequently decided to move the  $\lambda = 375$  nm laser to a dedicated fibre.



**Figure 4.20:** PCF focal shift measurements. (left) Waist radius vs position along the optical axis. (right) Fitted focal shift vs wavelength.

To reduce Rabi frequency drifts, it is useful to stabilise the beam polarization. We achieve this by aligning the dipole laser polarization to the fibre PM axes<sup>48</sup>. In order to achieve the optimal conditions for EIT cooling, we create  $\pi$ -polarized  $\lambda = 397$  nm light by aligning its fibre axis along the B-field. The repump lasers should then be coupled to the other PCF axis to produce  $(\sigma^+ + \sigma^-)$  light, which can repump all  $D_{3/2}$  and  $D_{5/2}$  sub-levels.

Alas, the original setup made no allowance for polarization control before the fibre input. While we managed to place an additional half-wave plate (HWP) in the  $\lambda = 866$  nm beam path and align its polarization correctly, this created a  $\approx 65\%$  loss on the (polarization-

<sup>43</sup> PCF aeroGUIDE-5-PM from NKT Photonics

<sup>44</sup> PLN 10X from Olympus

<sup>45</sup> AC254-200-A from Thorlabs

<sup>46</sup> 25 mm travel stages from OptoSigma

<sup>47</sup> RD-VL two-axis lens mount from Radiant Dyes, with micrometre screws


<sup>48</sup> It is not necessary for the PI lasers

dependent) dichroic beam combiner. It is fortunately not a big problem, since the saturation intensity of  $\lambda = 866$  nm is so low. On the other hand, we found no place to insert a polarizer in the  $\lambda = 854$  nm beam path, and that light is currently predominantly  $\pi$ -polarized. This significantly slows down the repumping rate out of the  $|D_{5/2}, m_J = \pm 5/2\rangle$  levels.

After many years of problem-free operation, we started noticing a recurrent degradation of the PCF tip (Fig. 4.21b) and the collimated beam profile (Fig. 4.21a). The contamination could be removed by polishing the tip, but would always re-appear, sometimes within a month, and sometimes within a few hours. We hypothesised that the  $\lambda = 397$  nm light charges the tip surface and attracts dust particles, but we do not know why the problem only affects this fibre<sup>49</sup> and why it started suddenly. Eventually, we managed to protect the spot from deterioration by constantly flowing nitrogen through the collimator tube [Mar17]. We later replaced the fibre with a new one, but are yet to test whether it can be stably operated without nitrogen.



(a) Degraded PCF spot as seen on a card inserted after the collimator.



(b) Degraded PCF tip seed under a microscope. After polishing, the fibre core becomes clean, though residue remains visible on the cladding.

**Figure 4.21:** PCF output degradation

### Photoionization beam delivery

The second-stage PI light source (initially  $\lambda = 375$  nm, later  $\lambda = 388$  nm) is delivered through a SM fibre and focused with a  $f = 200$  mm lens<sup>50</sup> onto the ion. The initial beam alignment is performed by the two mirrors, and the lens is manipulated using a two-axis translation mount<sup>51</sup> for fine beam pointing.

<sup>49</sup> The same PCF fibre is used across the lab in four different experiments, and this behaviour was never observed

<sup>50</sup> LA4102-UV from Thorlabs

<sup>51</sup> ST1XY-S/M from Thorlabs

#### 4. FREE-SPACE AND INTEGRATED OPTICS

##### $\lambda = 397\text{ nm}$ $\sigma$ laser delivery

The  $\lambda = 397\text{ nm}$   $\sigma$  laser is delivered through a PM fibre, aligned with two mirrors and focused with a  $f = 200\text{ mm}$  doublet lens. A Glan-Taylor cube<sup>52</sup> is used to linearly polarize the light; we place it after the pointing mirrors which may introduce polarization impurities. A HWP placed before the polarizer is used to maximise the power transmission, and a QWP placed after the polarizer creates circular polarization at the ion. An additional HWP placed after the Glan-Taylor was found to improve the final polarization purity, likely due to window birefringence.

##### $\lambda = 729\text{ nm}$ debug laser delivery

The  $\lambda = 729\text{ nm}$  “debug” laser is delivered through a PM fibre, aligned with two mirrors and focused with a  $f = 200\text{ mm}$  doublet lens onto the ion. This path uses the same viewport as the  $\lambda = 388\text{ nm}$  beam, and a mirror is mounted in a flip mount to select which path is active. A Glan-Taylor polarizer cleans the polarization out of the fibre, and a HWP is used to maximise the Rabi frequency of the transition of choice.

##### Ablation laser delivery

The ablation beam is produced by a pulsed ND:Yag laser<sup>53</sup> mounted on the experimental breadboard, originally set up by P. Zhou in his Master’s thesis [Zho17]. The beam is aligned with a pair of mirrors and focused with a  $f = 250\text{ mm}$  lens onto a target<sup>54</sup> placed  $\approx 20\text{ mm}$  from the trap chip. The lens transforms an input beam of  $w = 1.33(5)\text{ mm}$  onto a spot size of  $w = 330(10)\text{ }\mu\text{m}$ . This is approximately  $M^2 = 11$  times more than the spot expected for a Gaussian beam, which reflects the multi-mode nature of the input light. We confirmed with the manufacturer that this value of  $M^2$  is common for our source<sup>55</sup>. Initially, we used  $\lambda = 1064\text{ nm}$  pulses and later added a frequency-doubling module to generate pulses as  $\lambda = 532\text{ nm}$ . The optical components in the chain are coated for both wavelengths<sup>56</sup>. The pulses are  $5(3)\text{ ns}$  long, fire at a rate of up to  $15\text{ Hz}$ , and their power can be adjusted up to  $50\text{ mJ}$ . The beam power can be adjusted in two steps. First, a lever on the laser selects the total output power. I would normally adjust and then unscrew this knob, as even the tiniest touch changed the pulse power significantly. Then, the pulse power can be fine-tuned with a HWP and a polarizer.

---

<sup>52</sup> GT10-A from Thorlabs

<sup>53</sup> Minilite II from Continuum

<sup>54</sup>  $1 \times 1 \times 1\text{ cm}^3$  block cut from a calcium ingot, Alfa Aesar, polished with its flat face facing the trap

<sup>55</sup> Minilite II from Amplitude Laser

<sup>56</sup> Narrowband AR YAG coating from Thorlabs



### 4.5.2. Ion imaging

Ion imaging is performed using a custom Schwarzschild objective designed and characterised by Florian Leupold [Leu15]. The objective is designed for cryogenic operation with a working distance of 18 mm and produces a magnified image at a distance of 80 cm. The light can be directed into either a PMT<sup>57</sup>, or a camera<sup>58</sup>, or split 50/50 between the two. We use the PMT for day-to-day operations, and the camera is mainly used for alignment. A band-pass filter for 397 nm<sup>59</sup> ensures only the ion fluorescence light reaches the detectors. This filter can be removed to monitor neutral fluorescence. The PMT is placed behind a  $\phi = 1$  mm aperture in the image plane to remove laser scatter and improve the overall SNR. The objective magnification of  $M = 81$  allows two ions spaced by  $\approx 2 \mu\text{m}$  to be resolved on the camera. The optimal imaging performance is achieved within a  $\pm 100 \mu\text{m}$  field of view. The objective is mounted on a piezo stack with x/y/z travel range of 25 mm<sup>60</sup>. The stack is screwed onto the chamber and grounded. Furthermore, a fine gold mesh<sup>61</sup> shields the ion from any charges on the objective.

In the cryostat 1.0, the photon collection efficiency was measured at 0.7% at the PMT. While the objective had a high NA of 0.55, corresponding to 8.2% of the solid angle, the inner mirror blocked the photons within the NA of 0.31, reducing the total solid angle coverage to 5.7%. Furthermore, the use of metallic coatings meant that only 72% of the incident light was transmitted through the objective. Those features were originally included to allow for simultaneous imaging of  $^{40}\text{Ca}^+$  and  $^9\text{Be}^+$ , but the setup only ever operated with  $^{40}\text{Ca}^+$ . In order to improve the readout fidelity, we planned to replace the objective with a single AR-coated aspheric lens with NA of 0.5, which we expected to increase the light collection efficiency by up to 60%. The lens was designed and fabricated to provide a diffraction-limited spot size over  $\pm 40 \mu\text{m}$  from the centre<sup>62</sup>.

However, upon trapping in cryostat 2.0, a big surprise came: due to an increased chip reflectivity, the photon collection efficiency of the objective increased by approximately 50% (Sec. 5.3)! Therefore, the desired fidelity improvement was achieved, and it wasn't necessary to replace the imaging objective anymore. "Perfect is the enemy of good", says the old proverb, and for once I decided to follow it by postponing the imaging upgrade indefinitely. The only modification of the cryogenic imaging optics was to replace an uncoated imaging viewport with a AR-coated one, increasing the photon counts by 6%.

### 4.5.3. Imaging and laser alignment

The imaging system is aligned by illuminating the trap with a blue light-emitting diode (LED) through the imaging viewport. We begin by placing an additional  $f = 250$  mm lens

<sup>57</sup> PMT H10682-210 from Hamamatsu

<sup>58</sup> EMCCD iXon DU-897D-CSO-UVB from Andor

<sup>59</sup> FF01-390/1825

<sup>60</sup> ANP101/RES/LT/UHV from Attocube


<sup>61</sup> MG-17 from Precision Eforming

<sup>62</sup> Custom order from Thorlabs

#### 4. FREE-SPACE AND INTEGRATED OPTICS

before the camera and moving the imaging objective away from the trap. This allows us to see a zone of  $300\text{ }\mu\text{m} \times 300\text{ }\mu\text{m}$  (magnification  $M \approx 30$ ), as shown in Fig. 4.22a.

We then centre the objective on the desired trapping location, turn off the LED, and send a repumper laser beam across the chip. We raise<sup>63</sup> the beam towards the chip surface to introduce intentional scatter. With about  $100\text{ }\mu\text{W}$  of  $\lambda = 866\text{ nm}$  light, we can clearly see the beam at the trap (Fig. 4.22b), and we align it to the centre of the image.



**Figure 4.22:** Steps involved in aligning the imaging and the laser beams to the trapping zone.

We then remove the additional lens and align the imaging objective to the trapping zone with a magnification of  $M = 81$  (Fig. 4.22c). Finally, we repeat the beam-pointing to finely align the laser to the trapping location. The signal is typically much less clear at this stage, and the only alignment signature is the scatter from trap coupler openings (Fig. 4.22d).

The procedure is repeated for all the laser beams. Special care must be taken when

<sup>63</sup> Because the chip is upside down


aligning the PI light at  $\lambda = 388\text{ nm}$  since trap scattering leads to long-lasting charge accumulation. We found that  $20\text{ }\mu\text{W}$  of power is sufficient to see that beam on the camera, and we try to never exceed that so as to minimise the amount of charging of the trap chip. This alignment is further simplified by the fact that the  $\lambda = 388\text{ nm}$  light counter-propagates with the PCF output, which serves as a rough pointing guide.

Once the beams are visible on the chip, we aim to align them onto the ion with minimal tilt. The beam tilt with respect to the chip can be seen by monitoring the reflection angle on the other side of the vacuum system or inferred from the amount of scatter measured on the camera. All the beams are pointed with a slight tilt towards the chip, which might introduce more scatter, but ensures the beam hits the ion.


Once the beams are aligned and tilted, we move them by  $\approx 50\text{ }\mu\text{m}$  away from the trap surface. At the same time, we bring the objective  $\approx 30\text{ }\mu\text{m}$  away from the chip<sup>64</sup>. At this point, it is typically possible to trap and see an ion.

After trapping, beam pointing can be directly optimised by measuring the ion signal. First, the PCF and the imaging objective are aligned to maximise ion fluorescence. The  $\lambda = 397\text{ nm}$   $\sigma$  beam can be optimised by measuring the optical pumping rate (Sec. 5.2.1), and the debug quadrupole beam by maximising the Rabi frequency. Finally, the  $\lambda = 388\text{ nm}$  intensity can be optimised by measuring the AC stark shift it exerts on the optical qubit transitions ( $\approx 2\pi \times 1\text{ kHz}$  for  $100\text{ }\mu\text{W}$  of optical power).

#### 4.5.4. Ablation laser alignment



(a) Diode beam (small bright spot in the middle) focused onto the ablation target



(b) Low-power ablation beam (bright spot in the middle) on the target. The increased spot size in the image is partly due to larger power, and partly due to larger mode size ( $M^2 = 11$ )



(c) The oven wires (two vertical lines in the foreground) that are in the way and have to be avoided.

The ablation target features a slot to guide the laser alignment. In order to avoid hitting the trap with the powerful laser pulses, the alignment is performed cautiously as follows.

<sup>64</sup> Due to high trap reflectivity, a collect fluorescence emitted directly from the ion, as well as fluorescence reflected off the chip. While the ion is trapped  $\approx 50\text{ }\mu\text{m}$  away from the chip, the maximum fluorescence is collected with the objective focal point  $\approx 30\text{ }\mu\text{m}$  from the chip. However, at this position, the ion is not perfectly in focus

**Table 4.2:** Published experiments where trap operation is combined with optical integration.

Purpose	Component(s)	Wavelength(s)	Ion	Year	Citation
Detection	Optical fibre	280 nm	Mg <sup>+</sup>	2010	[Van10]
Detection	Micromirror	397 nm	Ca <sup>+</sup>	2011	[Mer11]
Detection	Photodiode	422 nm	Sr <sup>+</sup>	2012	[Elt12]
Detection	Diffractive mirror	370 nm	Yb <sup>+</sup>	2017	[Gha17]
Detection	SNSPD	313 nm	Be <sup>+</sup>	2020	[Tod20]
Addressing	Optical fibre	674 nm	Sr <sup>+</sup>	2011	[Kim11]
Addressing	Waveguides, grating couplers	674 nm	Sr <sup>+</sup>	2016	[Meh16]
Addressing	Optical fibres, waveguides, grating couplers	405 - 1092 nm	Sr <sup>+</sup>	2020	[Nif20]
Addressing	Optical fibres, waveguides, grating couplers	729 nm	Ca <sup>+</sup>	2020	[Meh20]
Addressing	Waveguides, grating couplers	435 nm	Yb <sup>+</sup>	2020	[Ivo20]

The pulsed laser output is first overlapped with a low-power CW laser diode beam<sup>65</sup>. The diode beam is then sent into the chamber, and we monitor its scatter with a camera aimed from the same direction<sup>66</sup>. We align the beam until we see the focused spot on the target, and no scatter from other areas (Fig. 4.23a). We then switch the pulsed laser on at low power, which allows us to see the laser pulses scattering off the target (Fig. 4.23b), and perform final aim adjustments (at full power it will be impossible to see the spot cleanly). In our case, we had to carefully avoid scattering off the oven wires (Fig. 4.23c).

## 4.6. Trap-integrated optics

Since laser interactions are a critical component of trapped ion quantum computing, scaling to large qubit numbers requires a scalable optical interface. In an ideal large-scale quantum processor, operations such as Rabi oscillations, cooling, state preparation and readout should be conducted in parallel and with low cross-talk across the chip. In [Kim07] it was first suggested that this could be achieved by trap-integrated optics. In a large processor, a combination of active and passive elements could be used to distribute and pulse the light, as well as to collect the photons and convert them to electrical signals.

Since then, multiple experiments realised proof-of-principle demonstrations of integrated optical control. The full list (to my knowledge) of published results on trap-integrated optics is shown in Tab. 4.2. The first experiment to integrate an optical element into an ion trap used an optical fibre to collect ion fluorescence [Van10]. Since then, further experiments showed how collection efficiency can be enhanced by integrating metal micromirrors [Mer11] and planar diffractive mirrors [Gha17] into the trap surface. The first integration of active

<sup>65</sup> CPS980S from Thorlabs

<sup>66</sup> The final mirror is a D-mirror, allowing the camera to easily see the trap

optical elements involved a PIN photodiode sandwiched into a semi-transparent trap made out of indium tin oxide (ITO)-coated quartz [Elt12]. Recently, a superconducting nanowire single-photon detector (SNSPD) was integrated into a planar trap and used to record ion fluorescence with high efficiency [Tod20].

The first experiment to address the ion through integrated optics used an optical fibre to drive the quadrupole transition in  $\text{Sr}^+$  [Kim11]. While the experiment observed quantum jumps, fibre-induced charging proved too strong to observe clean Rabi oscillations. The first coherent trap-integrated operations were demonstrated in [Meh16], where the light was delivered via integrated waveguides and focused directly onto the ion via diffractive grating couplers. However, the laser power efficiency was very low ( $-33$  dB input-to-ion), caused mainly by a free-space laser-chip interface. The trap used in this thesis is an upgraded version of that demonstration, with a direct fibre attach and a total loss of 6.4 dB (excluding fibre coupling). Simultaneous work demonstrated the integration of all wavelengths necessary for trapping and controlling  $\text{Sr}^+$ , albeit with a large loss of up to 35 dB [Nif20]. Most recently, an array of waveguides and grating couplers was integrated in a room-temperature  $\text{Yb}^+$  trap with 20  $\mu\text{m}$  ion-electrode distance, and coherent oscillations were demonstrated with 22 dB loss and a free-space input [Ivo20].


#### 4.6.1. Waveguides vs optical fibres

Optical waveguides and optical fibres are both composed of a core of refractive index  $n_1$  surrounded by cladding with lower refractive index  $n_2$ . However, a typical single-mode fibre operates with  $n_1 - n_2 \approx 0.01$ , while an optical waveguide features a strong index contrast. For example, the waveguides used in our trap feature a  $\text{Si}_3\text{N}_4$  core with  $n_1 = 2.0$  surrounded by a  $\text{SiO}_2$  cladding with  $n_2 = 1.45$ . The strong index mismatch increases the confinement strength, decreasing the necessary mode area and reducing bending losses. Waveguides are therefore a great way to route light on a chip, but they are not desirable for long-range transmission due to increased losses. The absorption and scattering are primarily caused by the roughness of the core-cladding interface [Lee00]. For example, the SM600 fibre we use to deliver the quadrupole light to the waveguide has a specified loss of  $\leq 15$  dB/km, while our on-chip waveguides have a measured loss of 2 dB/cm. While we expect the waveguide loss to improve with better fabrication [Bau10], it will remain orders of magnitude worse than that of an optical fibre.

#### 4.6.2. Waveguide modes


To understand the design of rectangular two-dimensional waveguides which route the light in our trap, it is helpful to examine the basic mode structure of a one-dimensional slab dielectric waveguide, as illustrated in Fig. 4.24. A slab waveguide supports two classes of modes propagating along  $\hat{z}$ , called transverse electric (TE) and transverse magnetic (TM). For TE modes, the electric field is oriented along  $\hat{y}$ , while for the TM modes, the magnetic (H) field is oriented along  $\hat{y}$ , and the electric field is primarily  $x$ -polarized. The

#### 4. FREE-SPACE AND INTEGRATED OPTICS



**Figure 4.24:** One-dimensional slab waveguide. It comprises of out a high-index core (refractive index  $n_1$  and thickness  $t$  along  $x$ ) sandwiched between a cladding of index  $n_2$ . The guided beam travels along  $z$ . The waveguide extends into infinity along the other transverse direction ( $y$ ).

mode structure can be obtained by finding the eigenvalues of the Helmholtz wave equation with appropriate boundary conditions. Fig. 4.25 shows the calculated fundamental modes for waveguides with varying waveguide thickness  $t$ , following the results in [Her84].



**Figure 4.25:** Analytically calculated modes of one-dimensional waveguides. These are solutions to the 1D Helmholtz equation with  $n_1 = 2.0$  for  $|x| < t/2$ ,  $n_2 = 1.45$  for  $|x| \geq t/2$  and  $\lambda = 729$  nm as described in [Her84]. The mode propagation direction is  $z$ , and the waveguide extends to infinity along  $y$ . (Left) The electric field of TE<sub>0</sub> modes for varying waveguide thickness  $t$ . For all the curves, the electric field is oriented along  $y$ . (Right) Comparison of the electric field magnitude of the TE<sub>0</sub> mode (red,  $y$ -polarized) and the TM<sub>0</sub> mode (blue,  $x$ -polarized) for core thickness  $t = 0.1$   $\mu\text{m}$  (shaded region). The majority of power is outside the core, and the TM<sub>0</sub> extends further into the cladding.

Fig. 4.25 (left) shows the effect of the waveguide thickness on the mode TE<sub>0</sub> mode size. For a very thick waveguide  $t \gg \lambda$ , the TE<sub>0</sub> is predominantly confined within the core. The overall size of the fundamental mode is comparable to the core thickness, and multiple orders of TE modes are allowed. As the mode thickness  $t$  decreases, the mode size follows suit, and the structure becomes single-mode. However, for a very thin waveguide ( $t \ll \lambda$ ), the fundamental mode becomes very weakly confined, and it expands again. Now, the

majority of power is outside the core, and the electric field decays exponentially with  $|x|$ . These observations hold for both the TE and the TM modes.

In order to create polarization-maintaining structures, we exploit an asymmetry between the TE and the TM modes illustrated in Fig. 4.25 (right). For the TE0 mode, the electric field (oriented along  $y$ ) is continuous at  $x = \pm t/2$ . This is however not the case for the TM0 mode, whose *displacement field* (oriented along  $x$ ) is continuous across the boundary, leading to:

$$n_1^2 E(x) \Big|_{|x| \rightarrow t/2^-} = n_2^2 E(x) \Big|_{|x| \rightarrow t/2^+}$$

where on the left-hand  $x$  side approaches  $t/2$  “from the inside” of the core and on the right-hand side  $x$  approaches  $d$  “from the outside” of the core. Since  $n_1 > n_2$ , near the boundary the electric field inside the core is lower than the electric field outside the core. This implies that the TM0 mode extends further along  $x$  than the TE0 mode, an effect which is particularly pronounced for  $t \ll \lambda$ . In Sec. 4.6.3, this effect will be exploited to attenuate the TM0 mode relative to the TE0 mode, thus creating a polarizing structure.

The actual trap waveguides are non-square rectangular waveguides, and the full treatment is significantly more cumbersome. Still, we distinguish between two fundamental modes (quasi-TE and quasi-TM) with predominant polarization along the two axes orthogonal to the propagation direction. Waveguide asymmetry lifts the mode degeneracy, and the modes do not couple to each other, creating a polarization-maintaining structure.


### 4.6.3. Trap with integrated waveguides

The trap, as well as all of its integrated optics, was designed by K. Mehta and fabricated by LioniX. All the information in this subsection is based on measurements conducted by K. Mehta and S. Miller. Further details on the measurements can be found in [Meh20] and [Mil19].

The trap features three trapping zones, each with two waveguide openings, as shown in Fig. 4.26. In each zone, one of the waveguides was designed for tight focusing of the  $\lambda = 729$  nm light, while the other produces a diverging beam for repumping (not used in this thesis).

#### Waveguide routing

The chip contains six light delivery waveguides and an additional looped waveguide for loss monitoring, for a total of eight fibre connections. Each waveguide comprises of a  $\approx 450$  nm  $\times$  200 nm  $\text{Si}_3\text{N}_4$  core, surrounded by  $\text{SiO}_2$  cladding. This structure is optimised for SM operation and low bending losses of the fundamental quasi-TE mode. Waveguide loss was measured with additional test structures containing waveguide loops of different lengths [Mil19], yielding a loss of 2 dB/cm. Further measurements with a tunable source



**Figure 4.26:** (left) Waveguide layout of the integrated-waveguide traps. The top and bottom waveguides travel in a loop and are used to align the fibre array and monitor the coupling losses. The remaining six waveguides travel to the three trap zones. (right) Microscope image of the trap centre, with trapping zones indicated by dashed circles. In every zone, the left waveguide is optimised for tight-focussing of the quadrupole laser, and the right waveguide is a simple grating to deliver defocused repump light (not discussed in this thesis). Source: K. Mehta.

revealed an increased attenuation at lower wavelengths with a scaling roughly consistent with the loss being dominated by sidewall roughness. This implies that a better fabrication procedures should allow to lower the loss.


### Fibre interface

To achieve an efficient fibre-waveguide interface, the mode exiting the optical fibre needs to match that of the waveguide. Since the waveguide provides much tighter confinement than a fibre, the waveguide mode needs to be expanded at the fibre interface. Two methods are commonly used for interfacing SM fibres with optical chips. The first is to connect the fibre vertically onto a grating coupler, which acts as a lens [Alo14]. The second method, which we chose to employ, is to use an edge coupler [Pap16]. In this design, the waveguide is adiabatically tapered as it approaches the chip edge. As discussed in Sec. 4.6.2, large modes can be created by either expanding or compressing the core. We chose to expand the waveguide along one direction while compressing it along the other, which creates a polarizing structure (see next subsection).

The details of our design are as follows. Eight flat-polished optical fibres<sup>67</sup> are held in a V-groove array with  $127\text{ }\mu\text{m}$  pitch. The eight waveguides are tapered by narrowing the waveguide height to  $25\text{ nm}$  and expanding its width from  $450\text{ nm}$  to  $700\text{ nm}$  over  $1\text{ mm}$ , as shown in Fig. 4.27 (note that the propagation direction is now  $x$ . We changed the notation as to make  $z$  the vertical direction, which matches the trap axis notation). As a result, the beam expands in both directions, and the total simulated mismatch between the fibre mode (MDF =  $5.4\text{ }\mu\text{m}$ ) and the waveguide mode is only  $-1\text{ dB}$ . Before coupling the fibre array to

<sup>67</sup> Nufern S630-HP





**Figure 4.27:** Waveguide cross section. Lighter colors indicate higher field intensity of the quasi-TE mode. (left) Thin waveguide cross-section for efficient fibre mode-matching. (left) Thick waveguide cross-section optimised for low-loss routing. Source: K. Mehta.

the chip, the trap edge needs to be polished. This is done to both reduce surface roughness, and to introduce a level of curvature onto the facet, which ensures a mechanical contact between the waveguide and the fibre, protecting it from subsequently introduced epoxy. The full procedure is described in detail in [Meh20], and proceeds in the following steps:

1. A 500  $\mu\text{m}$  glass piece is gold-coated and epoxied onto the trap chip
2. The trap chip is coated in a protective photoresist
3. The waveguide edge is polished on a fibre polishing paper
4. The Si substrate and the glass piece are grounded with silver epoxy
5. The chip and the fibre array are aligned and pressed together with translation stages and heated to 95  $^{\circ}\text{C}$ . The alignment is performed by monitoring transmission through the waveguide loop.
6. A small amount of epoxy is applied to the two edges of the interface and left to cure for two hours.


One of the big uncertainties of the project was the change in coupling loss between the room temperature and the cryogenic environment. To my surprise, we do not observe dramatic coupling changes between 300 K and 7 K. In the first sample (trap #3), we measured a relatively high loss of 5 dB per coupler at room temperature and in the cryostat. In the second sample (trap #4), the loss was measured to be 1.4 dB after attachment. Before starting the sample cool-down, we turned on the trap heater, and saw the loss increase to 2.4 dB. We then reduced the heater power, but the loss increase persisted, and remained at the same level throughout the cooldown and at 7 K. The sample was subsequently temperature-cycled three times. On the first two temperature cycles, the loss remained at 2.4 dB, while on the third one, it increased to 3.6 dB. In the final (trap #5), the loss was 3 dB at room temperature and increased to 3.9 dB upon the first cool-down, where it remained over the next multiple temperature cycles.

### Light polarization

The waveguide cross-section is not square, with a particularly large aspect ratio at the coupling region. The light entering the waveguide is not necessarily polarized, and couples into both the quasi-TE mode (polarized along  $y$ ) and the quasi-TM mode (polarized along  $z$ ). As discussed in Sec. 4.6.2, the narrow waveguide thickness causes the quasi-TM mode to extend further along  $z$ . The structure was designed to create a significant overlap between this mode and the Si substrate ( $2.7\text{ }\mu\text{m}$  away) and the metal ground plane ( $3.5\text{ }\mu\text{m}$  away). The overlap with non-dielectric materials causes significant loss, attenuating the quasi-TM mode to a negligible power after the 1 mm taper. The subsequent highly-confining waveguide supports both modes, but is polarization-maintaining, thanks to its non-rectangular cross-section. Thus, the light diffracting off the grating coupler is predominantly  $y$ -polarized.

### Light out-coupling

The light is coupled out of the chip with grating couplers. The couplers for  $\lambda = 729\text{ nm}$  light are based on a design principle presented in [Meh17]. The idea is two-fold. Along the propagation direction, the grating is patterned such that the light emitted towards the ion satisfies the Bragg condition. While the optimisation of the actual structure is done



**Figure 4.28:** Illustration of the Bragg diffraction on the grating couplers. The laser arrives through a rectangular waveguide from the left. The waveguide turns into a Bragg grating. The first-order diffraction is directed towards the ion, and the total diffracted power is increased by constructive interference of the upwards- and downwards-diffracted light. Note that the angle  $\theta$  in the figure corresponds to  $(-\theta)$  in the main text

with numerical tools, the following simple model allows us to understand the fundamental physics of the process. Consider a beam of free-space wavelength  $\lambda$  and wavevector  $k_0 = \frac{2\pi}{\lambda}$  travelling along a diffraction grating with period  $\Lambda$ , as shown in Fig. 4.28. The grating is composed out of two materials with refractive indices  $n_1$  and  $n_2 < n_1$ . In the limit of weak refractive index perturbation, and for a highly confined mode, we can consider the light as a monochromatic beam with a wavevector  $k = k_0 n_{\text{eff}}$ , where the effective refractive index satisfies  $n_2 > n_{\text{eff}} > n_1$ , depending on the grating duty cycle. Light that matches the Bragg condition is diffracted by the grating, and refracted on the cladding/air interface.

The Bragg condition for diffraction at an angle  $\theta$  to the vertical (in free space) reads


$$k_0 \sin \theta + k_0 n_{\text{eff}} = m\beta,$$

where  $m \in \mathbb{Z}$  and  $\beta = \frac{2\pi}{\Lambda}$ . The diffraction therefore occurs at

$$\theta = \arcsin \left( n_{\text{eff}} - \frac{m\beta}{k_0} \right).$$

This expression has a number of consequences. Since  $n_{\text{eff}} > 1$ , the  $m \leq 0$  emission is not present, and the first-order diffraction occurs for  $m = 1$ . By picking  $\beta > k_0 n_{\text{eff}}$ , can ensure that only one order undergoes constructive interference, and hence all emission occurs in one direction. We can then adjust the exact value of  $\theta$  by adjusting  $n_{\text{eff}}$  and  $\Lambda$ . The overall diffraction efficiency can be adjusted by varying the grating length, with longer gratings diffracting more power.

In our traps, we pick  $\Lambda = 300$  nm, which results in  $\frac{\beta}{k_0} = 2.43$  for  $\lambda = 729$  nm. Constructive interference occurs at  $-\theta = 36$  deg, corresponding to  $n_{\text{eff}} = 1.84$ . The grating spans a length of  $\approx 7.5$   $\mu\text{m}$ , which sets the beam waist radius at the ion (50  $\mu\text{m}$  above the chip) to  $w_x = 6.5$   $\mu\text{m}$ . This beam size was selected for efficiently simultaneous driving of two  $^{40}\text{Ca}^+$  ions spaced by 5  $\mu\text{m}$ . The total grating efficiency is simulated to be 50%, aided by constructive interference between the direct upwards reflection and the beam retro-reflected from the bottom silicon layer. Tight focusing along the transverse direction is




**Figure 4.29:** SEM image of a grating coupler optimised to focus light at  $\lambda = 729$  nm. Source: K. Mehta.

achieved by expanding the grating in a parabolic arc shown in the SEM image in Fig. 4.29. The grating curvature is set to create a  $w_y = 3.7$   $\mu\text{m}$  waist radius at the ion positions. Although it is possible to produce tighter spots, it is challenging to exactly align them to the pseudopotential null. The waveguide trap was specified (and measured) to have  $< 2$   $\mu\text{m}$  misalignment between the waveguide layer and a metal layer, and so setting  $w_y = 3.7$   $\mu\text{m}$  was a reasonable compromise. In addition, the grating lines in the three trapping zones

#### 4. FREE-SPACE AND INTEGRATED OPTICS

were deliberately displaced from the waveguide tapers, creating spots offset by  $(-2, 0, 2) \mu\text{m}$  in the  $y$  direction, such that at least one zone could be used in case of a larger misalignment, but this has proved to be unnecessary.

Prior to ion trapping, the optical emission profile was measured at the ion height with a 0.95 NA microscope, confirming the simulated spot sizes. The emission angle matches the design to within  $\approx 1$  deg, and the position of the focus is accurate to within  $2 \mu\text{m}$ , which is the optical measurement accuracy.




**Figure 4.30:** Simulated electric field distribution at the beam focus in the plane of the trap. (left) The dominant electric-field component is  $y$ -polarized (right) The dominant secondary polarization component ( $x$ -polarized). Note that the secondary polarization component originates predominantly from the grating coupler, i.e. it is present even when the input mode is pure quasi-TE. Source: K. Mehta.

The grating coupler largely preserves the dominant polarization direction of the incident quasi-TE mode, and the focused beam is predominantly polarized along  $y$ , i.e. the tight-focussing direction. However, away from the optical axis, secondary polarization components appear, predominantly along  $x$ , as shown in Fig. 4.30. While such a phenomenon is common to all tightly focused Gaussian beams, it is more pronounced for our grating coupler and depends on the particular taper geometry of the grating [Meh17]. For example, at  $y \approx \pm 2 \mu\text{m}$ , the magnitude of the  $x$ -component of the electric field peaks to  $E_x/E_y \approx 0.3$ . This is approximately three times larger than what is expected for a Gaussian beam [Tho12]. The  $z$ -polarized component peaks at a similar position, but with  $E_z/E_y \approx 0.1$ . This beam profile leads to possible highly non-trivial effects on Rabi frequencies, which is extensively discussed in the PhD thesis of C. Zhang [Zha21].

#### 4.6.4. Fibre routing

Fig. 4.31 shows how  $\lambda = 729 \text{ nm}$  light is routed from the AOM breadboard to the ion trap. As motivated in Sec. 4.1.4, the chain only consists of SM (non polarization-maintaining) fibres. The light is pulsed with a fibre-coupled single-pass AOM. After the AOM, an in-line polarization controller<sup>68</sup> is used to match the polarization into the waveguides. A 3 m-long SM optical fibre delivers the light to the experimental table, where 1% of power is picked-off

<sup>68</sup> CPC900 from Thorlabs



**Figure 4.31:** Integrated delivery of the  $\lambda = 729$  nm laser

with a fibre beamsplitter<sup>69</sup> for pulse power monitoring. The main output of the beamsplitter is connected to an 8-channel vacuum fibre feedthrough. In order to achieve high polarization stability, the long fibre section is placed in a plastic hose and taped to the optical table. Shorter fibre sections are carefully taped onto the optical breadboards whenever possible.

The vacuum feedthrough is a homemade 8-channel CF adapter. We make it by drilling through holes for fibre mating sleeves in a blank flange. Afterwards, eight flat-polished ferrule-terminated fibres are inserted from the air side into the mating sleeves, and the holes are sealed with Stycast. On the vacuum end, the feedthrough connection is made by inserting a fibre into the mating sleeve, while on the air side, we connect the input fibre through an APC/APC adapter to the feedthrough fibre. We were aware that fibre-to-fibre connectors can drift at cryogenic temperatures [Chr10; Maz18], and so all the fibre mating is done at room temperature. The eight fibres are then routed through openings in the heat shields and the 4K chamber and terminate with the trap V-groove array.

Throughout this thesis, we only used a single  $\lambda = 729$  nm channel at a time. At some points, additional  $\lambda = 780$  nm diode light was used to induce trap substrate photoconductivity (Sec. 8.2.1) and for trap charge stabilisation (Sec. 8.2.2). This light was either passed through the waveguide loop (to scatter into the silicon substrate) or combined with the quadrupole light (to stabilise the total power at the grating coupler). This light is not used anymore and hence excluded from the diagram.

<sup>69</sup> TW670R1A1 from Thorlabs



## 5. Single-ion operations with integrated waveguides

This section discusses the calibration and benchmarking of single-ion operations in the integrated waveguides traps. Sec. 5.1 describes the experimental pre-requisites: ion loading and micromotion compensation. We then discuss different methods of state preparation (Sec. 5.2) and single-ion readout (Sec. 5.3). The discussion of coherent operations begins with measurements of qubit lifetime and coherence (Sec. 5.4). That section also introduces the formalism used to quantify experimental noise, which is used to benchmark single-qubit rotation errors in Sec. 5.5. Finally, Sec. 5.6 turns the attention to the motional state, and we discuss the calibration of motional state cooling, as well as heating rate measurements and their interpretation.

### 5.1. First steps

#### 5.1.1. Loading ions

To load an ion, we first generate neutral atoms either from a resistive oven or from an ablation source (Sec. 3.2.5). The presence of neutral atoms crossing the trap can be observed by neutral fluorescence induced by the  $\lambda = 423\text{ nm}$  laser. At that point, we turn on the photoionization laser, and continuously perform a Doppler cooling sequence (Sec. 2.5.2) followed by qubit readout (Sec. 2.3.6). We also add an additional strong ( $> 10\text{ }\mu\text{W}$ ) tone on the  $\lambda = 397\text{ nm}$   $\pi$  laser which is far red-detuned ( $-2\pi \times 160\text{ MHz}$ ) from the detection resonance [vMou21]. We block the neutral fluorescence light at  $\lambda = 423\text{ nm}$  with a filter and wait to observe fluorescence at  $\lambda = 397\text{ nm}$ , which signifies a trapping event. Immediately after trapping, the neutral atom source and the neutral fluorescence laser and the photoionization laser are switched off.

In the setup 1.0, neutral atoms were always generated by an effusive oven located at the 4K stage, which caused a number of issues (Sec. 3.2). In the setup 2.0, we initially attempted to solve those problems using laser ablation, which is discussed in Sec. 5.1.3. We subsequently returned to using an effusive oven, this time located at the room-temperature stage, and attempted to improve the loading rates by changing the photoionization wavelength from  $\lambda = 375\text{ nm}$  to  $\lambda = 388\text{ nm}$ , as described in Sec. 5.1.4. While loading is an ongoing struggle, the room-temperature oven clearly outperforms the cryogenic one, and during the latest opening, we removed the latter source from the 4K chamber altogether.

## 5. SINGLE-ION OPERATIONS WITH INTEGRATED WAVEGUIDES

In a SET, trapping is only possible in a narrow parameter range. Specifically, we need to apply precise compensation fields along the radial directions to approximately align the DC null with the RF pseudopotential minimum. The correct fields can be found by systematic micromotion compensation (Sec. 5.1.2). This makes it challenging to trap for the first time in any given trap, or to re-trap following a strong substrate charging event. In that scenario, we blindly scan the offset fields until an ion is trapped, and then proceed with micromotion compensation. Assuming all lasers are correctly aligned, this typically takes no longer than a few hours in a new trap.

### 5.1.2. Micromotion minimisation

After the first successful trapping in a new trap, it is necessary to minimise the micromotion. We do this by applying additional fields along the radial directions, translating the position of the DC null along  $z$  (vertical, perpendicular to the chip surface) and  $y$  (horizontal, parallel to the chip). The position  $x$  of the DC null along the trap axis does not participate in the micromotion compensation, and can be chosen freely to maximise the Rabi frequency of the trap-integrated light. However, different positions  $x$  correspond to different amounts of axial micromotion.


We proceed as follows. First, the  $y$  and  $z$  offset fields are adjusted to increase the total fluorescence. At this point, optical qubit carrier transition is usually clearly identifiable with Rabi spectroscopy, but the laser-ion interactions are still strongly modulated by the micromotion. This is due to the Doppler shift caused by ion oscillation. For ion micromotion described by time-dependent displacement  $\vec{a} \sin(\Omega_{\text{RF}} t)$ , the coupling terms in the interaction Hamiltonian (Eq. (2.6)) become modified according to:

$$e^{-i\omega t} \rightarrow e^{-i(\omega t + \beta \sin(\Omega_{\text{RF}} t))} = e^{-i\omega t} \sum_{n=-\infty}^{\infty} J_n(\beta) e^{in\Omega_{\text{RF}} t},$$

where  $\beta = \vec{k} \cdot \vec{a}$  is known as the *modulation index*, and the  $J_n(x)$  are the Bessel functions of the first kind. This expression shows micromotion modulation sidebands, which are detuned by  $\pm n\Omega_{\text{RF}}$  from the carrier transitions. In order to minimise  $\beta$ , we scan the  $y$  and  $z$  offset fields until the Rabi frequency of the first micromotion sideband (proportional to  $J_1(\beta)$ ) is minimised [Kel15]. This requires use of two different beam directions. The trap-integrated light is sensitive to vertical micromotion, while the free-space “debug” beamline resolves the horizontal micromotion. Fig. 5.1 shows the typical time-traces of the first micromotion sideband as a function of  $y$  and  $z$  offset fields. The micromotion-compensated position in Fig. 5.1 is  $z \approx 14.2 \mu\text{m}$  and  $y \approx 1.98 \mu\text{m}$ , and corresponds to the point at which the sideband Rabi frequencies are minimised.

One subtlety of this method is that  $\beta = 0$  is not the unique minimum of  $J_1(\beta)$ , and it is possible to extinguish the first micromotion sideband without eliminating micromotion. To exclude this possibility, we use the fact that  $\beta = 0$  is the only simultaneous minimum of





**Figure 5.1:** Rabi oscillations on the first micromotion sideband vs compensation field along  $y$  (parallel) and  $z$  (normal to the trap surface). The displacement magnitude is calculated assuming mode frequency  $\omega_{y,z} = 2\pi \times 3$  MHz. In reality, modes are higher frequency, tilted and non-degenerate, therefore a compensation field along  $z$  does not displace the ion purely along the trap normal. (left) The micromotion sideband driven with the trap-integrated laser (36 deg to the trap normal). (right) The micromotion sideband driven with the debug laser (parallel to the chip, 45 deg to the trap axis)

both  $J_1(\beta)$  and  $J_2(\beta)$ . Therefore, we repeat the measurements on the second micromotion sideband and verify that its Rabi frequency is also locally minimised.

In all the waveguide-integrated traps, the necessary offset field is predominantly vertical, implying that the excess electric field is dominated by the trap itself. We typically adjust the  $z$  compensation on a daily basis, while drifts in  $y$  are rare and smaller in magnitude. While the initial micromotion calibration is challenging and time-consuming, the daily calibration only requires a single repetition of the scan in Fig. 5.1 (left) and is very rapid.

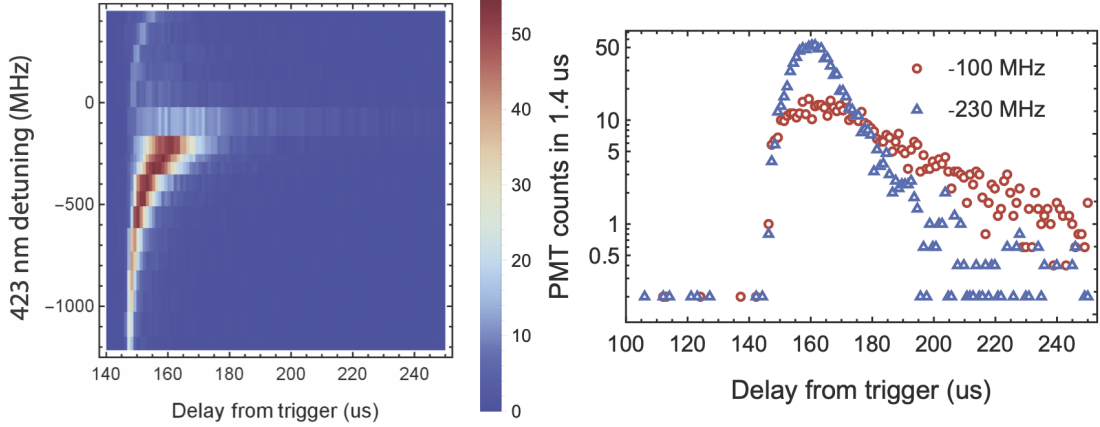
We also employed other micromotion compensation techniques. The RF amplitude modulation technique (a.k.a. “RF tickling”) [Iba11] was used in traps without integrated light delivery, as well as to benchmark axial micromotion. In addition, we investigated the possibility of minimising micromotion in all directions with the single trap-integrated beam. This relied on the fact that while transverse micromotion causes no Doppler shift, it leads to amplitude-modulation when the laser is tightly focused. These methods and results are described in detail in the PhD thesis of C. Zhang [Zha21].

### 5.1.3. Ablation loading

In order to alleviate the problem caused by running an oven in the 4K chamber, we attempted to load using laser ablation. We began the loading attempts by measuring neutral fluorescence generated by 2 mJ pulses at  $\lambda = 1064$  nm. The fluorescence was measured by opening the PMT for  $t = 1.4$   $\mu$ s at a variable delay time from the pulse trigger. Fig. 5.2 (left) shows the neutral fluorescence signal as a function of detuning  $\delta$  of the

## 5. SINGLE-ION OPERATIONS WITH INTEGRATED WAVEGUIDES

$\lambda = 423$  nm laser from the atomic resonance. Following [Gug15] we can fit the observed Doppler shift, finding that the neutral plume makes an average angle of  $\theta \approx 6.5$  deg to the photoionization beam. While we observe the largest atom numbers at  $\delta \gtrsim 2\pi \times 200$  MHz, these are predominantly fast-moving and unlikely to become trapped. Instead, we tuned the  $\lambda = 423$  nm frequency to  $\delta \approx 0$  where we observe the long tail of slow-moving atoms (Fig. 5.2, right). Following the calibration, we began the loading attempts. At the pulse



**Figure 5.2:** (left) The neutral fluorescence signal generated by a single ablation pulse in a 1.4  $\mu$ s time window. The maximum signal is  $\approx 50$  counts, where the PMT saturates. As the  $\lambda = 423$  nm frequency is red-detuned from the natural frequency, the peak fluorescence corresponds to a shorter delay from the trigger, i.e. faster-moving atoms. This is caused by a Doppler shift consistent with an angle of  $\theta \approx 6.5$  deg between the atoms and the laser. The minimum delay time of 145  $\mu$ s is almost all due to the delay between the electronic trigger and the pulse generation time. The frequency is measured with 100 MHz resolution. (right) Comparison of the signals at  $\delta = -100$  MHz and  $\delta = -230$  MHz. While the latter trace shows a larger total signal (area under the curve), the former trace exhibits a longer tail of slow-moving atoms. We find that the loading probability is the largest closer to the natural frequency.

energy of 1 – 2 mJ, ablation loading was possible, but slow: at a repetition rate of 15 Hz, loading would take several minutes, corresponding to thousands of shots. Increasing the pulse energy to  $> 5$  mJ increased the loading rate significantly, and an ion would often appear after a few shots. However, at this energy, we observed significant stray field drifts, making loading unreliable.

We then frequency-doubled the pulsed laser to a wavelength of  $\lambda = 532$  nm, where high-efficiency loading was reported in  $^9\text{Be}^+$  [Leo19]. We tested the effect of the ablation power on the stray electric fields by trapping an ion and monitoring its fluorescence and position as we applied ablation pulses of increasing power. For pulse energies below 700  $\mu$ J, we observed little to no drifts of the trapping fields. However, starting at  $\approx 700$   $\mu$ J, the effect suddenly became dramatic, as every 20 pulses caused a visible ion displacement of  $\approx 3$   $\mu$ m along the trap axis. This happened even in absence of photoionization light. We attempted to load ions at pulse energy  $< 700$   $\mu$ J, but not a single ion was trapped even for

tens of thousands of repetitions.

All of these findings have a clear interpretation. In order to trap with high efficiency, the ablation plume needs to contain a lot of neutral atoms, which is achieved for high ablation pulse energies. However, at those high energies, the ablation plume also contains ions, which become implanted on the trap surface, affecting the stray electric fields. At low pulse powers, the per-pulse loading probability is too low for a low repetition-rate laser to be useful. For example, the red points in Fig. 5.2 (right) show a total of 580 neutral counts per pulse, or 9000 counts per second for a 15 Hz repetition rate. This is an order of magnitude less than the typical neutral fluorescence signal of our atomic oven. In all likelihood, other experiments which use low repetition-rate ablation pulses are either less sensitive to trap charging (e.g. no exposed dielectrics) [Lei07] or have a significantly higher per-atom ionization probability.


The obvious solution to this problem is to use a high repetition-rate pulsed laser [Hen07; Pog21]. This should allow for rapid loading without direct ion production. Note that for this purpose it is actually beneficial to keep the focused pulse spot large to avoid crossing the ionization threshold. Alternatively, better dielectric shielding should make us less sensitive to charge deposition. Finally, the per-ion loading probability could be enhanced by larger photoionization laser powers.

#### 5.1.4. Atomic oven loading

Following the ablation investigations, we decided to focus on optimising the oven loading. First, we installed an additional oven at a room temperature stage (Sec. 3.1.5). This allowed us to run the oven for longer periods of time with little vacuum degradation or 4K chamber heating.

Our next attempt was to improve the per-ion loading rate by increasing the photoionization power. Specifically, we measured that the initial  $\lambda = 375$  nm power ( $\approx 50 \mu\text{W}$ ) was far from saturation, and that increasing it significantly enhanced the loading rate. However, those observations were difficult to implement in practice, since that laser caused significant trap charging. An alternative suggestion was to replace the  $\lambda = 375$  nm source with a one nearer the photoionization threshold, with some theoretical calculations implying this could lead to a  $100\times$  increase of the scattering cross-section [Lec10]. However, photoionization with a  $\lambda = 388$  nm source diode did not lead to a significant change of loading rates one way or another, as shown in Fig. 5.3 (left). We later discovered that more thorough theoretical calculations do not in fact predict a dramatic wavelength sensitivity of the cross-section [vdHar03].

Right now, the loading rates are only changing slowly from day to day, but vary significantly from month to month. One consistent observation is that the micromotion-compensated *working-well* often is not the voltage set with the highest loading probability. Instead, we often manage to find a voltage set at the same axial location which exhibits consistently higher loading rates, as shown in Fig. 5.3 (right). Those *loading-wells* are



**Figure 5.3:** (left) Loading rates for different second-stage photoionization powers. The measurements were performed a few months apart in different traps and at  $\approx 120 \mu\text{W}$  of  $\lambda = 423 \text{ nm}$  power. While increasing the second-stage PI power increases the loading rate, moving the wavelength closer to the photoionization threshold does not. The measurements show the loading rates to be identical between the different traps. That is a coincidence, as the loading rates drift by a factor of 2-3 from month to month. Nonetheless, the wavelength change does not dramatically affect the ion loading. (right) Comparison between the loading rates in the micromotion-compensated working-well and a dedicated loading-well. The loading-well consistently outperforms the working-well on the same day and at the same  $\lambda = 423 \text{ nm}$  power.

found by trial-and-error and do not persist in the long term. It would be very beneficial to understand how to systematically construct and calibrate loading-wells. However, we have thus far not studied which parameters of the loading wells could account for their success, and whether pushing the DC null closer to further away from the chip is beneficial. Loading-well calibration is also the weak point of two-ion loading (Sec. 8.1.1)

## 5.2. State preparation


All single-ion experiments begin by preparing the ion into a well-defined state. We selected  $|S_{1/2}, m_J = -1/2\rangle$  as the target state of the optical pumping, which is the  $|1\rangle$  qubit state of both the optical and the Zeeman qubit (Sec. 2.2.7). Sec. 5.2.1 describes the calibration and performance of the dipole transition optical pumping step, which prepares  $|1\rangle$  with a probability of  $\approx 0.999$ . Faced with this limitation, we tuned the quadrupole transition optical pumping sequence as described in Sec. 5.2.2, reducing the error to  $4(1) \times 10^{-5}$ .

### 5.2.1. Dipole transition optical pumping

#### Beam alignment


The fidelity of dipole transition optical pumping depends on the degree of selectivity of the  $\Delta m = -1$  transition, which requires the pumping beam to be perfectly circularly polarized and perfectly aligned along the quantization axis. The  $\lambda = 397 \text{ nm}$   $\sigma$  beam is aligned in steps. First, the beam is linearly polarized and aligned onto the ion by maximising fluorescence. We then insert a QWP and rotate it to roughly maximise the contrast of Rabi oscillations on the  $|S_{1/2}, m_J = -1/2\rangle \rightarrow |D_{5/2}, m_J = -1/2\rangle$  transition. We then insert an additional HWP to compensate for the viewport birefringence.

The final precise alignment is performed by walking the beam pointing, as well as the angles of both the HWP and the QWP. The first step of the precise alignment is to run



**Figure 5.4:** Pulse sequence for dipole transition optical pumping calibration. The ion is prepared in  $S_{1/2}$  after Doppler cooling. The  $\lambda = 397 \text{ nm}$   $\sigma$  beam is aligned to minimise the probability to drive an ion out of  $|S_{1/2}, m_J = +1/2\rangle$  with a carrier pulse. A short  $\lambda = 866 \text{ nm}$  pulse is inserted together and after the  $\sigma$  light to avoid population trapping in  $D_{3/2}$ .


the optical pumping followed by a probe pi-pulse on a spectator transition  $|S_{1/2}, m_J = +1/2\rangle \rightarrow |D_{5/2}, m_J = -3/2\rangle$ , shown in Fig. 5.5. The probe pulse is made deliberately slow to avoid off-resonant excitation of any transition out of  $|S_{1/2}, m_J = -1/2\rangle$ . The beam pointing and polarization are adjusted until the excitation probability is minimised. At



**Figure 5.5:** Pulse sequence for fine calibration of dipole transition optical pumping. A long  $\lambda = 397 \text{ nm}$   $\sigma$  is applied with the  $\lambda = 866 \text{ nm}$  light switched off. This causes population pumping into  $D_{3/2}$ , which is then detected as a bright ion during the readout step. The  $|S_{1/2}, m_J = -1/2\rangle$  state is shelved into  $D_{5/2}$  by a pair of  $\lambda = 729 \text{ nm}$  pulses, and appears dark during the detection.

this point, further alignment is challenging as the real-time signal is very small, making it challenging to distinguish between waveplate settings a few degrees apart. To that end, we have developed a second method to improve the polarization alignment (Fig. 5.5). The state is (imperfectly) initialised in  $|S_{1/2}, m_J = -1/2\rangle$  by dipole transition optical pumping. We subsequently apply  $\lambda = 397 \text{ nm}$   $\sigma$  pulse of duration  $\tau$ . Due to polarization impurity, this pulse causes residual scattering from  $|S_{1/2}, m_J = -1/2\rangle$ , pumping the population into the meta-stable state  $|D_{3/2}\rangle$  with a constant rate. Afterwards, we execute the Zeeman

qubit readout sequence (Sec. 5.3.2), shelving the  $|S_{1/2}, m_J = -1/2\rangle$  population into the  $|D_{5/2}\rangle$  manifold. Subsequent “bright” measurement registers the combined population of  $|S_{1/2}, m_J = +1/2\rangle$  and  $|D_{3/2}\rangle$ . The power and detuning of the scattering beam are adjusted until a clear repumping signal can be observed. We then rotate the waveplates until the scattering rate reaches a minimum. Fig. 5.6 shows how this method allows us to see a difference between the optimal QWP angle and a 3 deg angle offset.



**Figure 5.6:** Example traces of the fine dipole transition optical pumping calibration. As is clear, even small errors in waveplate orientation can be resolved. The signal corresponds to population pumping to  $D_{3/2}$ , which was independently verified.


### Performance and limit

Since spin preparation requires only a few photons to be scattered, we always have sufficient power to saturate the process within  $t = 1.4\mu\text{s}$  (the minimal pulse time in our control system), and we have not performed a detailed process optimisation. By monitoring the depth of the optical Rabi oscillations from  $|S_{1/2}, m_J = +1/2\rangle$ , we find:

$$P(+1/2) \approx 1.0 \times 10^{-3} \quad \text{state preparation error, dipole transition optical pumping.}$$

Our best interpretation of this limit is that the process saturates due to effective polarization impurity created by the AC Zeeman shift from the trap RF current, illustrated in Fig. 5.7 [Cho10; Gan18]. The oscillating electrode voltage, combined with finite electrode impedance, produces a B-field component  $\vec{B}_0 = |B_0|\hat{y}$  aligned at 45 deg to the quantization field  $\vec{B} = |B|(\hat{x} + \hat{y})/\sqrt{2}$ , modulating its strength and direction. We estimate the value of  $|B_0|$  in two ways. First, we simulate the B-field produced by current  $I$  on the RF electrode. The current magnitude  $|I|$  at the trap centre can be estimated from:

$$|I| = \frac{1}{2}\Omega_{\text{RF}}V_{\text{RF}}C,$$



**Figure 5.7:** Directions of the dominant B-field components. The quantization axis (blue) is oriented at 45 deg to the trap axis ( $x$ ). The RF current flowing along  $x$  (red) creates a B-field component along  $y$  (green), modulating the overall B-field strength and orientation at the RF frequency.

where  $\Omega_{\text{RF}} = 2\pi \times 35 \text{ MHz}$  is the RF frequency,  $V_{\text{RF}} = 30 \text{ V}$  is the RF amplitude, and  $C = 10 \text{ pF}$  is the trap capacitance. The factor of  $1/2$  reflects the fact that the current in the centre of the trap is approximately 50% of the current that flows through the wirebonds (the other end of the trap imposes a boundary condition of  $I = 0$ ). This consideration results in  $|I| = 33 \text{ mA}$ , which produces  $|B_0| = 0.56 \text{ G}$  in the simulation. The value of  $|B_0|$  can be also measured by driving the micromotion sidebands of the Zeeman qubit transition. This method and is described in the thesis of C. Zhang and leads to a very precise estimate of  $|B_0| = 0.6014(19) \text{ G}$ , proving that the RF electrode current is the dominant cause of the residual B-field fluctuations [Zha21].

We can now estimate the state preparation error as follows. Suppose that the  $\lambda = 397 \text{ nm}$   $\sigma$  light is perfectly aligned along  $\vec{B}$  and is perfectly circularly polarized. However, the total B-field direction oscillates with a time-dependent angle  $\theta(t)$ . In the limit of  $|B_0| \ll |B|$ , this is given by:

$$\theta(t) = \frac{1}{\sqrt{2}} \frac{|B_0|}{|B|} \sin(\Omega_{\text{RF}} t).$$

For a small static angle error  $\theta$ , the  $\Delta m = -1$  transition is no longer the only allowed excitation even for ideal polarization, and the first-order error comes from the  $\Delta m = 0$  transition, whose matrix element is  $\propto \theta^2/2$ . In our case, the angle is not static, since the B-field oscillation frequency ( $\Omega_{\text{RF}} = 2\pi \times 35 \text{ MHz}$ ) is comparable to the atomic excitation rate ( $\Gamma = 2\pi \times 21.6 \text{ MHz}$ ). Nonetheless, we will continue the calculation as an order-of-magnitude estimate. In the first order, the steady state is obtained by equalising the rates  $\Gamma_\pi, \Gamma_\sigma$  of the two dominant dissipative processes

$$\begin{aligned} \Gamma_\pi &= \Gamma \left( |S_{1/2}, m_J = -1/2\rangle \rightarrow |P_{1/2}, m_J = -1/2\rangle \rightarrow |S_{1/2}, m_J = +1/2\rangle \right) \\ \Gamma_\sigma &= \Gamma \left( |S_{1/2}, m_J = +1/2\rangle \rightarrow |P_{1/2}, m_J = -1/2\rangle \rightarrow |S_{1/2}, m_J = -1/2\rangle \right) \end{aligned}$$

which occur at rates:

$$\begin{aligned}\Gamma_\pi &\propto s(\delta_\pi) \times P(-1/2) \times \frac{\langle \theta^2 \rangle}{2} \times \frac{2}{3} \\ \Gamma_\sigma &\propto s(\delta_\sigma) \times P(+1/2) \times \left(1 - \frac{\langle \theta^2 \rangle}{2}\right) \times \frac{1}{3},\end{aligned}$$

where  $s(\delta)$  is the lineshape function (Sec. 2.3.4), and  $\delta_{\pi/\sigma}$  are the detunings from the respective transitions. The last factor of each equation is the Clebsch-Gordan coefficient of each decay channel (App. 11.1). The equilibrium condition evaluates to:


$$P(+1/2) = \frac{s(\delta_\pi)}{s(\delta_\sigma)} \langle \theta^2 \rangle + O(\langle \theta^4 \rangle)$$

Setting  $s(\delta_\pi)/s(\delta_\sigma) \approx 0.5$  and  $\langle \theta^2 \rangle = B_0^2/(4B)$  gives the state-preparation error of

$$P(+1/2) \approx \frac{1}{8} \left( \frac{B_0}{B} \right)^2 \approx 10^{-3}.$$

While the calculation is only approximate, it strongly hints that the measured state-preparation error could be dominated by the B-field oscillations. Other possible error sources include beam polarization impurity or a beam pointing error<sup>1</sup>.

### 5.2.2. Quadrupole transtion optical pumping



**Figure 5.8:** Pulse sequence used to calibrate quadrupole transition optical pumping. The method consists of  $N$  loops of a  $\lambda = 729$  nm pi-pulse from  $|S_{1/2}, m_J = +1/2\rangle$  followed by a repump to  $S_{1/2}$ . The pumping error is quantified by measuring the contrast of a Rabi oscillation on the same transition.


The state preparation error can be improved by applying quadrupole transition optical pumping after the dipole transition optical pumping. We use the scheme described in Sec. 2.3.5 using the  $|S_{1/2}, m_J = +1/2\rangle \rightarrow |D_{5/2}, m_J = -3/2\rangle$  drive. This transition was chosen for two reasons. First, it has the lowest frequency out of all optical qubit transitions, reducing spectral crowding. Second, the dominant<sup>2</sup> dissipation channel  $|D_{5/2}, m_J = -3/2\rangle \rightarrow |P_{3/2}, m_J = -3/2\rangle$  favours the decay to  $|S_{1/2}, m_J = -1/2\rangle$  over the

<sup>1</sup> In this case, an external B-field leads to a strong static tilt of the quantization, which requires a large tilt of the beam angle to compensate. While in principle correctable, there might be local optima of the polarization/pointing alignment that make the global optimum hard to find

<sup>2</sup> Recall that the  $\lambda = 854$  nm laser is predominantly  $\pi$ -polarized, thus favouring the  $\Delta m = 0$  transition



other ground state, speeding up the convergence. Following the optical pumping, we drive a pi-pulse on the  $|S_{1/2}, m_J = +1/2\rangle \rightarrow |D_{5/2}, m_J = -3/2\rangle$  transition to benchmark the error, as shown in Fig. 5.8.



**Figure 5.9:** Results of the quadrupole transition optical pumping calibration. The probability of exciting a transition out of  $|S_{1/2}, m_J = +1/2\rangle$  decreases exponentially with the number of loops, saturating at  $6(1) \times 10^{-5}$ . The limit is a combination of the state preparation error, as well as a small probability for the shelving pulse to excite transitions out of  $|S_{1/2}, m_J = -1/2\rangle$ .

Fig. 5.9 shows the measured error versus the number loops. The error decreases exponentially with the number of loops, saturating at

$$P(+1/2) = 6(1) \times 10^{-5}.$$

A single loop takes  $\approx 12 \mu\text{s}$ . The measured error cannot be accounted for by the readout histogram overlap error (Sec. 5.3). We verified spectroscopically that off-resonant couplings to spectator transitions should likewise be negligible, which was confirmed by measuring  $P(+1/2)$  to be unchanged when the quadrupole Rabi frequency was decreased by  $\times 2$ . The most likely explanation for the observed limit is the laser spectral impurity (Sec. 4.4.8). Direct measurements reveal that the spin-polarization pulses excite the population out of  $|S_{1/2}, m_J = -1/2\rangle$  with probability  $p = 2.0(5) \times 10^{-5}$  per cycle, in qualitative agreement with the spectral impurity data in Fig. 4.18. Those measurements suggest that the actual state preparation error is

$$P(+1/2) = 4(1) \times 10^{-5}$$

with the remainder attributable to the readout imperfection. If correct, this suggests that the state preparation error can be lowered by up to a factor of 5 by increasing the B-field magnitude, beyond which the laser spectrum needs to be narrowed.

A practical challenge of robust quadrupole transition optical pumping in our experiment is the polarization of the  $\lambda = 854 \text{ nm}$  laser. As explained in Sec. 4.5, this laser is unfortunately

close to  $\pi$ -polarized, suppressing the repumping out of the  $|D_{5/2}, m_J = \pm 5/2\rangle$  states. In every loop, the state can decay from  $|P_{3/2}, m_J = 3/2\rangle$  to  $|D_{5/2}, m_J = 5/2\rangle$ , where it can become trapped. This manifests itself as a spin polarization convergence slow-down. The problem can be mitigated by increasing the  $\lambda = 854$  nm laser power, rotating its polarization or adding a short magnetic dipole rotation pulse between the  $D_{5/2}$  sub-levels.

### 5.3. State readout

#### 5.3.1. Optical qubit readout

Optical qubit readout (Sec. 2.3.6) is calibrated by fitting experimental histogram counts with two Poisson distributions with means  $\mu_D$  and  $\mu_B$  corresponding to the dark ( $|0\rangle = |D_{5/2}\rangle$ ) and bright ( $|1\rangle = |S_{1/2}\rangle$ ) states respectively. We typically use  $\mu_D = 1.7, \mu_B = 48$  for a detection time  $t_D = 250$   $\mu$ s. At the same time, we can estimate the readout saturation ( $n_P$  in Eq. (2.23)) by fitting the fluorescence lineshapes obtained by scanning the laser cavities with steady-state solutions of the optical Bloch equations for a three-level system [Leu15]. The means above correspond to  $n_P \approx 0.18$ , implying an overall detection efficiency of  $\eta \approx 0.8\%$ . This is  $\approx 60\%$  more than the detection efficiency estimated with the same method in setup 1.0, and we attribute the increase to the high reflectivity of the trap chip.

The dominant contribution to the detection error with the threshold method is the finite optical qubit lifetime. Intuitively, the finite lifetime  $\tau = 1.1$  s of  $|0\rangle$  leads to a non-zero probability that an initially dark ion begins fluorescing during the detection, and appears bright overall. Formally, we model this process by adding a non-Poissonian component to the overall histogram  $H(n) = D(n) + S(n)$  (see [Bur10] for details):

$$D(n) = e^{-t_D/\tau} \frac{\mu_D^n e^{-\mu_D}}{n!} + X(n)$$


$$S(n) = \frac{\mu_B^n e^{-\mu_B}}{n!}$$

where

$$X_0(n) = \int_{\mu_D}^{\mu_D + \mu_B} g(\lambda) \frac{\lambda^n e^{-\lambda}}{n!} d\lambda,$$

$$g(\lambda) = \frac{1}{\mu_B} \exp\left(\frac{\lambda - \mu_B - \mu_D}{\mu_B}\right).$$

The simulated fluorescence histogram is shown in Fig. 5.10 (left). The detection error  $\epsilon_0$  ( $\epsilon_1$ ) is defined as the probability that  $|0\rangle$  ( $|1\rangle$ ) is classified as bright (dark). For a given



**Figure 5.10:** (left) Simulated histogram for an ion prepared in  $|0\rangle$  (red) or  $|1\rangle$  (blue). The yellow trace shows the histogram for  $|0\rangle$  for a qubit with infinite lifetime. (right) The resulting readout errors as defined in the main text. Due to finite qubit lifetime, the  $|0\rangle$  readout error  $\epsilon_0$  increases to  $> 1 \times 10^{-4}$  for the optimal threshold, whereas the  $|1\rangle$  readout error  $\epsilon_1$  remains negligible.

threshold  $n_{\text{th}}$ , those error can be calculated from:

$$\begin{aligned}\epsilon_0 &= 1 - \text{CDF}(D(n_{\text{th}})) \\ \epsilon_1 &= \text{CDF}(S(n_{\text{th}}))\end{aligned}\tag{5.1}$$

where  $\text{CDF}(F(m)) = \sum_{n=0}^m F(n)$  is the cumulative distribution function of the distribution  $F$ . The simulated detection errors are shown in Fig. 5.10 (right).

In the absence of decay, both detection errors are negligible. The decay however introduces a significant error asymmetry  $\epsilon_0 \gg \epsilon_1$  around the optimal threshold. For our parameters and  $n_{\text{th}} = 14$  we find:

$$\epsilon_0 = 1.7 \times 10^{-4} \quad \text{and} \quad \epsilon_1 = 2.2 \times 10^{-9}$$


The first error can be decreased by decreasing the detection time  $t_D$ , allowing for reduction to  $\epsilon_0 \approx 9 \times 10^{-5}$  at  $t_D = 130 \mu\text{s}$ . Below  $t_D$ , the histogram overlap error begins to increase. On the other hand, high-precision readout can still be performed for certain input states by exploiting the low value of  $\epsilon_1$ . Consider a general qubit input state  $|\psi\rangle = \cos(\theta/2)|0\rangle + e^{i\phi}\sin(\theta/2)|1\rangle$  with ideal measurement outcomes  $P(0) = \cos^2(\theta/2)$  and  $P(1) = \sin^2(\theta/2)$ . Our non-ideal threshold measurement records [Hom06]

$$\begin{pmatrix} P(D) \\ P(B) \end{pmatrix} = \begin{pmatrix} 1 - \epsilon_0 & \epsilon_1 \\ \epsilon_0 & 1 - \epsilon_1 \end{pmatrix} \begin{pmatrix} P(0) \\ P(1) \end{pmatrix}.\tag{5.2}$$

The overall measurement error  $\epsilon_\psi$  can be defined as:

$$\epsilon_\psi = |P(B) - P(1)|$$

which is plotted in Fig. 5.11. As is clear, high-fidelity measurements can be performed for



**Figure 5.11:** Overall measurement error vs qubit angle  $\theta$  on the Bloch sphere. States in the vicinity of  $\theta = \pi/2$  (with a large overlap with  $|1\rangle$ ) are measured with very low error due to low decay probability.

$\theta \approx \pi/2$ . We use this method to benchmark the quadrupole transition optical pumping error (Sec. 5.2.2) with a readout error  $< 10^{-6}$ . We also note that while the thresholding technique is the simplest one, detection errors can be improved in real-time by Bayesian adaptive maximum-likelihood estimation [Mar20] or in post-processing by directly fitting the measurement histograms [Neg18a]. However, we have not use either of these techniques in this thesis.

### 5.3.2. Zeeman qubit readout

As discussed in Sec. 2.3.7, the Zeeman qubit can be read out by performing a number of optical pi-pulses, shelving one of the Zeeman states into  $D_{5/2}$ , followed by an optical qubit readout. The state transfer probability of a single optical pi-pulse is  $\approx 0.999$  (Sec. 5.5.1). By performing two shelving pulses on two different transitions, the shelving error should decrease to  $\approx 10^{-6}$ , and thus the Zeeman qubit can be read-out with essentially the same fidelity as the optical qubit. The choice of shelving transitions is somewhat arbitrary, but care must be taken that the shelving pulse does not excite any other transition off-resonantly. At the moment, the shelving pulse always excites some transitions off-resonantly due to laser spectral impurity, limiting the readout error to  $\gtrsim 10^{-5}$  (Sec.5.2.2), although we have not characterised this error carefully.

Since carrier Rabi frequencies decrease with ion temperature (Sec. 5.6.1), the duration of the shelving pulses should likewise be adjusted to the motional state. This makes it challenging to perform experiments with long and variable wait times. The main effect of ion heating is a decrease of shelving probability, and we can model this error through Eq. (5.2) with  $\epsilon_1 = 0$ , leading to:

$$P(B) = \epsilon_0 P(0) + P(1)$$


In order to estimate  $\epsilon_0$  in any given experiment, we perform it twice, once as desired, and once with a Zeeman pi-pulse inserted right before detection. That inverts the qubit populations, and the second experiment registers a bright ion with probability

$$P'(B) = \epsilon_0 P(1) + P(0),$$

and we can therefore estimate  $\epsilon_0$  from

$$\epsilon_0 = P(B) + P'(B) - 1.$$

Whenever we need to perform a Zeeman readout after a long wait time, we first measure  $\epsilon_0$  while scanning the duration and number of the shelving pulses until  $\epsilon_0$  is minimised. Fig. 5.12 shows the residual growth of  $\epsilon_0$  for the Zeeman qubit Ramsey experiment with spin-echo (Fig. 5.20). By appropriate shelving pulse choice, we keep  $\epsilon_0 < 0.5$  for wait



**Figure 5.12:** Shelving error  $\epsilon_0$  for the Zeeman qubit Ramsey experiment with spin-echo. The linear trendline shows a gradual increase of the shelving error as the motional temperature increases. The oscillations that deviate from the line are an artefact of the insufficient averaging: the oscillation frequency recorded in the experiment drifts slowly over minutes, and the two datasets necessary to reconstruct  $\epsilon_0$  were taken with a few minutes of delay.

times up to 150 ms, where the axial mode temperature  $\bar{n} \approx 300$ . Additionally, whenever we measure significant  $\epsilon_0$ , we correct for the bias it introduces while post-processing the results. This is achieved by inverting the matrix in Eq. (5.2), resulting in:

$$\begin{pmatrix} P(0) \\ P(1) \end{pmatrix} = \begin{pmatrix} 1/(1 - \epsilon_0) & 0 \\ -\epsilon_0/(1 - \epsilon_0) & 1 \end{pmatrix} \begin{pmatrix} P(D) \\ P(B) \end{pmatrix}.$$

## 5.4. Dephasing and decay

Before benchmarking the quality of single-qubit rotations, we turn our attention to the simplest gate of all: the identity gate. The qubit error introduced by simply waiting sets the baseline for all other operations, which are almost certain to experience more noise. We


distinguish two types of errors: qubit decay and qubit dephasing. Qubit decay is associated with a population transfer between the qubit states  $|0\rangle$  and  $|1\rangle$ . In the Markovian limit, the decay timescale is usually referred to as  $T_1$ . Qubit decoherence quantifies the phase error introduced by waiting. Coherence loss is caused by qubit frequency changes in the frame of the drive; thus, it is produced both by atomic level fluctuations and by drive frequency fluctuations. Qubit decoherence can be measured with a Ramsey sequence, and its timescale is usually referred to as  $T_2^*$ . This dephasing is usually caused by a combination of Markovian and non-Markovian processes, and the latter can be eliminated by suitable dynamical decoupling sequences. Coherent error suppression can increase the decay timescale up to a limit referred to as  $T_2$ .

This section is organised as follows. Sec. 5.4.1 describes the observation of enhanced optical qubit decay, and how it was reduced to the natural lifetime of  $T_1 = 1.1$  s. Sec. 5.4.2 introduces the Ramsey sequence and the dynamical decoupling methods used to estimate qubit dephasing rate and its origin. Following that, we describe the decoherence measurements on the optical qubit (Sec. 5.4.3) and the Zeeman qubit (Sec. 5.4.4).

#### 5.4.1. Optical qubit lifetime

All optical qubit operations are limited by the natural decay rate of the  $D_{5/2}$  state  $\tau = 1.1$  s. The finite-lifetime error is typically a negligible correction compared to other coherent operation errors, although it does limit the readout fidelity (Sec. 5.3.1). However, the first measurement of the  $D_{5/2}$  decay rate revealed a much shorter lifetime of  $\approx 50$  ms.

The origin of the problem was the leakage of nominally extinguished light at  $\lambda = 854$  nm when the double-pass AOM was switched off. The issue was two-fold. First, we realised we were operating the laser with powers vastly exceeding the saturation intensity. This allowed us to obtain robust repumping in  $t < 1.4$   $\mu$ s, and made us insensitive to the PDH cavity drifts. While convenient, this mode is susceptible to leakage errors. By placing the cavity on resonance, we not only increase the desired repumping rate for a given intensity, but also attenuate the repumping rate of the leakage light, which has a different frequency. In our experiment, the leakage light is detuned by  $\delta = 2\pi \times 160$  MHz from the repump light, resulting in significant attenuation on a transition with  $\Gamma = 2\pi \times 1.35$  MHz. The second problem was that the low extinction ratio of the double-pass AOM setup. We improved it by moving the beam block closer to the beam focus, and then further by tilting it with respect to the optical axis. We benchmark the performance of the new system by preparing the optical qubit in  $|D_{5/2}, m_J = -1/2\rangle$  and measuring its decay to  $S_{1/2}$ . We superimpose the observed decay curves with simulations of the 3-level system dynamics to extract the Rabi frequency of the  $\lambda = 854$  nm drive. The results are shown in Fig. 5.13. We reduced the overall power until the repumping occurs at a timescale comparable with the minimum experimental pulse time of  $t = 1.4$   $\mu$ s. The fitted Rabi frequency of the  $\lambda = 854$  nm laser is  $\Omega_{\text{on}} = 2\pi \times 1.8$  MHz. When the repumping light is extinguished, the optical qubit decay rate is consistent with the natural linewidth of the quadrupole transition. From the fits,




**Figure 5.13:** Measurements of the repumping rate from  $|D_{5/2}, m_J = -1/2\rangle$  to  $S_{1/2}$ , together with manual fits of three-level system dynamics with variable  $\lambda = 854$  nm Rabi frequency. (left) The repumping rate with  $\lambda = 854$  nm turned on is consistent with  $\Omega = 2\pi \times 1.8$  MHz. (right) With the  $\lambda = 854$  nm laser turned off, the measured decay is consistent with  $\Omega = 0$  (dashed line). The measurements can be used to approximately lower-bound the Rabi frequency to  $\Omega < 2\pi \times 10$  kHz (solid line).

we can upper-bound the Rabi frequency of the leakage light to  $\Omega_{\text{off}} < 2\pi \times 10$  kHz. This corresponds to an extinction ratio of:

$$\frac{P_{\text{on}}}{P_{\text{off}}} = \left( \frac{\Omega_{\text{on}}}{\Omega_{\text{off}}} \right)^2 > 3 \times 10^4.$$

If the extinction ratio is improved by approximately two orders of magnitude, the system will become insensitive to cavity drifts once again. This should be easily achievable by better double-pass AOM layout (App. 11.4).

#### 5.4.2. Ramsey experiment



**Figure 5.14:** A Ramsey sequence. A qubit prepared in  $|1\rangle$  is initialised in an equal superposition of  $|0\rangle$  and  $|1\rangle$  with a  $\pi/2$ -pulse. Following a free-evolution period of length  $\tau$ , the qubit phase is measured by a variable-phase  $\pi/2$ -pulse.

A standard tool for coherence measurement is a Ramsey experiment (Fig. 5.14) [Ram50]. A qubit is prepared in  $|1\rangle$  and a  $R(\pi/2, 0)$  pulse rotates it to  $(|0\rangle + i|1\rangle)/\sqrt{2}$ . A subsequent

## 5. SINGLE-ION OPERATIONS WITH INTEGRATED WAVEGUIDES

period of free evolution of duration  $\tau$  rotates the state to:

$$|\psi(\tau)\rangle = \frac{1}{\sqrt{2}} \left( |0\rangle + ie^{-i\phi(\tau)} |1\rangle \right),$$

where

$$\phi(\tau) = \int_0^\tau \delta(t) dt$$

is the integral of the time-dependent detuning  $\delta(t)$  in the frame of the drive. Following a subsequent  $R(\pi/2, \phi_2)$  rotation and a projective readout, we measure the probability  $P(1)$  for the qubit to return to  $|1\rangle$ . It is easy to show that:

$$P(1) = \frac{1}{2} (1 - \cos(\phi(\tau) - \phi_2)).$$

In the simplest form, Ramsey experiment can be used for qubit frequency calibration, since when the qubit frequency stays constant ( $\delta(t) = \delta_0$ ), the measurement outcome  $P(1)$  is constant with  $\tau$  if and only if  $\delta_0 = 0$ .

However, as the interrogation length  $\tau$  increases, the assumption of constant detuning eventually breaks down. The resulting qubit frequency noise results in phase noise on  $\phi(\tau)$ , limiting the achievable contrast. The resulting oscillations take the form of

$$P(1) = \frac{C(\tau)}{2} (1 - \cos(\phi(\tau) - \phi_2)).$$

with the contrast  $C(\tau)$  limited by the fluctuation in  $\phi(\tau)$ . Suppose that  $\delta(t)$  oscillates around  $\delta_0 = 0$ . Then  $\langle \phi(\tau) \rangle = 0$  and the contrast loss is specified only by the variance of  $\phi(\tau)$  [Sep19]:

$$C(\tau) = e^{-\frac{1}{2} \langle \phi^2(\tau) \rangle}.$$

The Ramsey contrast  $C(\tau)$  thus holds information about temporal correlations of  $\delta(t)$ . Suppose that  $\delta(t)$  is stationary, and hence can be fully described by an autocorrelation function  $R(\tau) = \langle \delta(t)\delta(t+\tau) \rangle$ . Then:

$$\langle \phi^2(\tau) \rangle = \int_0^\tau dt \int_0^\tau dt' R(t - t').$$

Intuitively, the autocorrelation function describes the duration and strength of noise correlations. At  $\tau = 0$  we find  $R(0) = \langle \delta^2(t) \rangle \equiv \sigma^2$ , which is the noise variance. On the other hand, for realistic noise sources, we expect  $R(\tau) \rightarrow 0$  as  $\tau \rightarrow \infty$  [Rie03].

In order to evaluate  $\langle \phi^2(\tau) \rangle$  for different classes of noise, it is useful to move to the Fourier



space, and describe the noise correlation in terms of its power spectral density  $S(\omega)$ <sup>3</sup>. This is accomplished by setting:


$$R(t - t') = \int_0^\infty e^{i\omega(t-t')} S(\omega) \frac{d\omega}{2\pi}. \quad (5.3)$$

The expression for phase variance can be now integrated to give:

$$\langle \phi^2(\tau) \rangle = \tau^2 \int_0^\infty S(\omega) \left[ \frac{\sin(\omega\tau/2)}{\omega\tau/2} \right]^2 \frac{d\omega}{2\pi}.$$

The quantity in the square brackets is called the *window function*. Intuitively, its shape selects which spectral components of the noise affect the final result, and which get attenuated. Here, the window function is peaked around  $\omega = 0$ , and thus low-frequency noise causes significant contrast loss.

### Spin echo



**Figure 5.15:** A Ramsey sequence with a spin echo. The period of free-evolution is divided into two segments of length  $\tau/2$ , and a  $\pi$ -pulse inserted in the middle is used to cancel common-mode errors.

In order to reduce the influence of the ubiquitous low-frequency noise, it is common to modify the Ramsey sequence to include a spin echo (also known as a Hahn echo) [Hah50]. This is accomplished by dividing the free-evolution period  $\tau$  into two sections of length  $\tau/2$  with a  $\pi$ -pulse in between (Fig. 5.15). The state at time  $\tau$  is then given by:

$$|\tilde{\psi}(\tau)\rangle = \frac{1}{\sqrt{2}} \left( |0\rangle + ie^{-i\tilde{\phi}(\tau)} |1\rangle \right),$$

where

$$\tilde{\phi}(\tau) = \int_0^{\tau/2} \delta(t) dt - \int_{\tau/2}^{\tau} \delta(t) dt, \quad (5.4)$$

and all subsequent relations can be derived analogously. For constant frequency  $\delta = \delta_0$  we find  $\tilde{\phi}(\tau) = 0$ , illustrating the insensitivity to low-frequency noise (frequency offsets and

<sup>3</sup> Although  $S(\omega)$  has units of frequency, it is not the laser lineshape!

## 5. SINGLE-ION OPERATIONS WITH INTEGRATED WAVEGUIDES

drifts). For stationary noise with mean 0, the spin echo contrast decays as:

$$\tilde{C}(\tau) = e^{-\frac{1}{2}\langle\tilde{\phi}^2(\tau)\rangle},$$

where

$$\langle\tilde{\phi}^2(\tau)\rangle = \tau^2 \int_0^\infty S(\omega) \left[ \frac{\sin^2(\omega\tau/4)}{\omega\tau/4} \right]^2 \frac{d\omega}{2\pi}. \quad (5.5)$$

Note in particular that the window function is now zero at  $\omega = 0$  and that the function periodicity changed by a factor of 2.

### Decay profiles

While the above expressions allow us to convert spectral density into phase variance, in reality, the noise spectral density is rarely known with high precision. In order to make progress, we will examine a number of simplified noise models and their influence on the Ramsey decay [OMa15]. The first limit is when the noise is high bandwidth, and so the noise spectral density is approximately constant within the window function. In this *white noise limit* we set  $S(\omega) = S_0$  in Eq. (5.4) and Eq. (5.5), which results in:

$$\langle\phi_{\text{white}}^2(\tau)\rangle = \langle\tilde{\phi}_{\text{white}}^2(\tau)\rangle = \frac{S_0}{2}\tau.$$

Therefore, in the white noise limit, the Ramsey experiment dephases exponentially with time, and at the same rate regardless of the echo:

$$C_{\text{white}}(\tau) = \tilde{C}_{\text{white}}(\tau) = e^{-\tau/T_{\phi_1}} \quad \text{where} \quad T_{\phi_1} = 4/S_0.$$

Note that this dephasing happens exponentially regardless of the exact lineshape outside of the window function. Intuitively, it does not matter how fast the noise is fluctuating, as long as it is fluctuating much faster than  $1/\tau$ .

In the laboratory frame, we can view qubit dephasing as caused by either by fluctuations of the transition frequency (e.g. due to B-field variation) or fluctuations of the oscillator frequency. In our experiments, the former is the dominant cause of Zeeman qubit dephasing (Sec. 5.4.4), while the latter is responsible for majority of the optical qubit dephasing (Sec. 5.4.3). In any case, oscillator linewidth sets an upper limit on the coherence time. Consider an oscillator whose frequency is described by a Lorentzian lineshape with FWHM angular linewidth  $\Delta\omega$ . This corresponds to a single-sided noise spectral density of

$$S_0 = \frac{2\Delta\omega}{\pi}.$$

Thus the relation between the linewidth and the coherence decay time reads:

$$\Delta\omega = \frac{2\pi}{T_{\phi_1}}$$

I emphasise that this expression is valid for white noise only, which may or may not accurately describe the phase noise of the drive.

The second limit is that of slow noise with long correlation times. In this limit, we can treat the noise as at a specific frequency  $\omega_0$ . Most commonly, the slow fluctuations will be peaked at  $\omega_0 = 0$ , modelled as  $S(\omega) = S_0\delta(\omega)$ , where  $\delta(\omega)$  is the Dirac delta function. Evaluating the integrals in Eq. (5.4) and Eq. (5.5) gives<sup>4</sup>:

$$\langle\phi_{\text{slow}}^2(\tau)\rangle = \frac{S_0}{4\pi}\tau^2 \quad \text{and} \quad \langle\tilde{\phi}_{\text{slow}}^2(\tau)\rangle = 0.$$

The resulting Ramsey contrast decays as:

$$C_{\text{slow}}(\tau) = e^{-(\tau/T_{\phi_2})^2} \quad \text{where} \quad T_{\phi_2} = \sqrt{\frac{8\pi}{S_0}},$$

$$\tilde{C}_{\text{slow}}(\tau) = 1.$$

Therefore, slow noise results in a Gaussian decay of the contrast, which gets completely suppressed by the spin echo.

While in the slow noise limit we do not think of the drive lineshape, we can relate  $S_0$  to the variance of the detuning. Using Eq. (5.3) and recalling  $R(0) = \sigma^2$  we find that:

$$S_0 = 4\pi\sigma^2.$$

Therefore, the relation between the decay time and the qubit frequency standard deviation reads:

$$\sigma = \frac{\sqrt{2}}{T_{\phi_2}}.$$

This calculation can be generalised to noise peaked at a specific frequency  $\omega_0$  by setting  $S(\omega) = \frac{1}{2}S_0\delta(\omega - \omega_0) + \frac{1}{2}S_0\delta(\omega + \omega_0)$ , where  $S_0 = 2\pi A^2$  for oscillation amplitude  $A$ . The usual integration gives:

$$\langle\phi_{\omega_0}^2(\tau)\rangle = \frac{S_0}{\pi\omega_0^2}\sin^2(\omega_0\tau/2) \quad \text{and} \quad \langle\tilde{\phi}_{\omega_0}^2(\tau)\rangle = \frac{4S_0}{\pi\omega_0^2}\sin^4(\omega_0\tau/4),$$

---

<sup>4</sup> Note that the integral goes from 0 to  $\infty$ , and so delta function integrates to 1/2


with a decay profile of:

$$C_{\omega_0}(\tau) = e^{-(A^2/\omega_0^2) \sin^2(\omega\tau/2)}$$

$$\tilde{C}_{\omega_0}(\tau) = e^{-(4A^2/\omega_0^2) \sin^4(\omega\tau/4)}.$$

The three possible decay profiles described in this section are by no means the only possibilities, and complicated lineshapes result in many different possible decay spectra. For example, 1/f noise (not covered by the theory above) causes an approximately Gaussian contrast decay regardless of the echo [Rie03]. Still, the tools described in this section allow us to begin understanding the coherence loss observed in the experiment. By fitting and comparing the decay curves of Ramsey fringes with and without the spin-echo echo, we will be able to say to what extent the decay is dominated by fast versus slow noise, and describe the spectral characteristics of the slow components.


### 5.4.3. Optical qubit Ramsey



**Figure 5.16:** Optical qubit Ramsey experiment contrast vs wait time  $\tau$ . The resulting fits include a single-frequency components as well as an overall decay term, which is either Gaussian (solid line) or exponential (dashed line). (left) Over the long timescales, it is difficult to separate the exponential model from the Gaussian one (right) At short timescales, the Gaussian model is clearly favoured by the data. Note that the observed decay for  $\tau < 100 \mu\text{s}$  is solely due to the single-frequency component in the Gaussian model.

Fig. 5.16 shows the results of a Ramsey experiment performed on the optical qubit transition  $|S_{1/2}, m_J = -1/2\rangle \rightarrow |D_{5/2}, m_J = -1/2\rangle$ , together with corresponding fits. The contrast decay exhibits a strong monochromatic component with frequency  $\omega_0 \approx 2\pi \times 175 \text{ Hz}$  and amplitude  $A \approx 2\pi \times 185 \text{ Hz}$ . The superimposed decay is most likely Gaussian, as the exponential fit significantly overestimates the contrast loss at short times  $\tau < 100 \mu\text{s}$ . The displayed Gaussian decay time is  $T_{\phi_2} \approx 11 \text{ ms}$ , corresponding to qubit frequency standard deviation of  $\sigma \approx 2\pi \times 20 \text{ Hz}$ . Note that the observed decay for  $\tau < 100 \mu\text{s}$  is caused solely by the monochromatic component, and the Gaussian decay is negligible. These observations also allow us to constrain the exponential decay timescale to  $T_{\phi_1} > 50 \text{ ms}$ .

(linewidth  $\Delta\omega < 2\pi \times 20$  Hz). The optical qubit spin-echo sequence, shown in Fig. 5.17, decays on a timescale of  $\tilde{T} \approx 20$  ms, and the strong monochromatic component makes it challenging to study the finer details of the spectrum.




**Figure 5.17:** Optical qubit Ramsey experiment with spin echo contrast vs wait time  $\tau$ . The fits include a single-frequency component with frequency and amplitude fixed to the value obtained from the regular Ramsey experiment.

The optical qubit decay is dominated by laser frequency noise and not by B-field noise. This can be inferred from the Ramsey experiments on the Zeeman qubit presented in the subsequent section 5.4.4, from which B-field noise can be extracted. Those results set the minimum dephasing rate of the optical qubit which would be expected due to B-field fluctuations alone to  $T \approx 65$  ms (no spin echo) and  $\tilde{T} \approx 1.5$  s (with spin-echo). Direct measurement of other optical qubit transitions (with higher B-field sensitivities) result in comparable decay times, confirming the hypothesis. We do not currently know what causes the monochromatic noise component, but optical beat-note measurements with an independent laser reveal a strong noise component at a comparable frequency which is present on all the  $\lambda = 729$  nm light locked to the high-finesse cavity (Sec. 4.4.8). Thus, it is unlikely the noise originates in the power amplification setup or due to the use of optical fibres and integrated optics. I think it may be an artefact of the PDH lock, or caused by acoustic noise on the cavity.

Approximately a year after the data presented above was taken, I carefully repeated the measurements in light of difficulties obtaining high-fidelity gates. The result, shown in Fig. 5.18, was a decay curve that does not agree with the previous one. While the monochromatic component decreased in amplitude to  $A \approx 2\pi \times 120$  Hz (at the same frequency), the short-time measurements revealed a *faster* decay, which can only be accounted by fast noise. While a pure exponential fit still overestimates the observed decay, a combined fit with  $T_{\phi_1} \approx 35$  ms and  $T_{\phi_2} \approx 11$  ms reproduces the data well.

The change is significant yet mysterious. The simplest guess at the possible origin of the


## 5. SINGLE-ION OPERATIONS WITH INTEGRATED WAVEGUIDES



**Figure 5.18:** Optical qubit Ramsey experiment contrast vs wait time  $\tau$  one year after Fig. 5.16. (left) The single-frequency component is decreased in amplitude compared to the data in Fig. 5.16 (right) Despite that, the short-timescale contrast loss is enhanced, and can no longer be fully accounted for by overall Gaussian decay (solid line) or exponential decay (dashed line). The data can be explained by a combination of a shorter-timescale Gaussian decay and a slower exponential decay (dotted line).


increased fast noise is the change in the cavity lock settings. Specifically, the FALC unit gain is adjusted from time to time while debugging lock problems, and it is challenging to adjust it reproducibly as it's controlled by a very touchy potentiometer. It is unclear whether the linewidth increase could be in any way related to the drop in acoustic noise, or whether it was coincidental. The data clearly indicates this issue requires more attention in the future, and experiments should examine to what extent the laser lock settings influence the noise spectrum.

### 5.4.4. Zeeman qubit Ramsey



**Figure 5.19:** Zeeman qubit Ramsey experiment contrast vs wait time  $\tau$ . Neither the Gaussian nor the exponential fit capture the decay or revival of the data very well.

Fig. 5.19 shows the contrast decay of the Zeeman qubit coherence. The decay is clearly more Gaussian than exponential, although neither curve captures the behaviour faithfully. The displayed Gaussian decay timescale is  $T_{\phi_2} = 12$  ms, corresponding to  $\sigma = 2\pi \times 19$  Hz. Since the Zeeman qubit splitting is directly proportional to the applied magnetic field, we can extract the standard deviation of B-field fluctuations of  $\sigma_B \approx 4 \mu\text{G}$ . This is only an approximation since the fluctuations seem to be actually caused by a range of slow-oscillating components. The presence of a discrete noise component is evidenced by the spin-echo curve



**Figure 5.20:** Zeeman qubit Ramsey experiment with spin echo contrast vs wait time  $\tau$ . The fits include a single-frequency component as well as an overall decay term.

displayed in Fig. 5.20. The echo data exhibits a significantly longer coherence, implying that the Ramsey decay is dominated by slow drifts. The dominant noise component has a frequency of  $\omega_0 = 2\pi \times 46$  Hz and an amplitude of  $A \approx 2\pi \times 43$  Hz. The remaining decay could be exponential ( $T_{\phi_1} = 400$  ms) or Gaussian ( $T_{\phi_2} = 200$  ms).

While the noise in the data does not allow for precise determination of  $\omega_0$ , the value is significantly shifted from the 50 Hz noise frequency. We also found that line-triggering the experiment does not affect the decay time, further strengthening the hypothesis that the dominant noise is not from mains. The measured coherence times significantly exceed those reported for room-temperature setups without magnetic shielding [Rus16; Lin20]. This suggests that the cryogenic chamber screens the high-frequency noise, though we cannot claim to have measured the shielding factor.

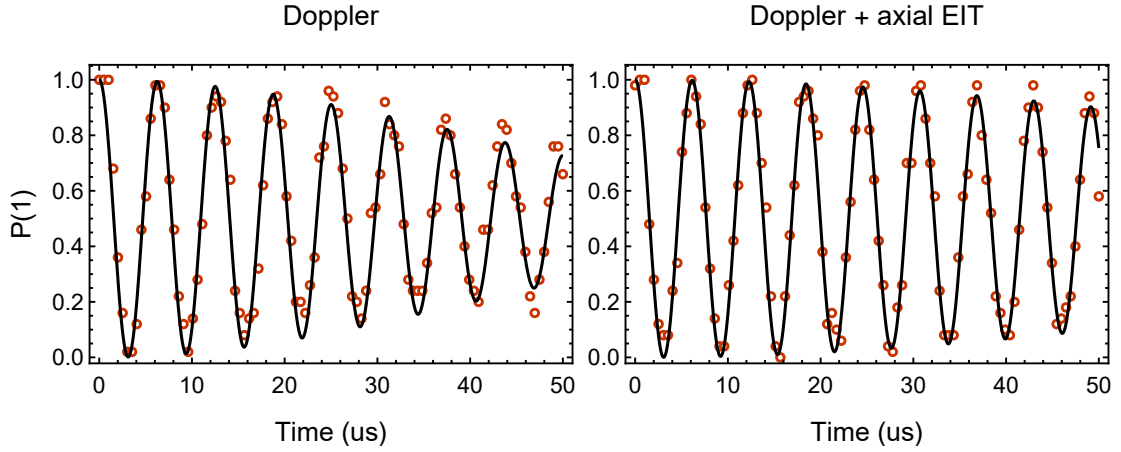
My best guess is that the residual noise component originates from differential vibrations between the magnets and the trap. Assuming a B-field gradient of  $\approx 5$  G/m due to the coil field inhomogeneity (Sec. 3.4.3 and Sec. 9.3.4), the observed noise could be accounted for by 100 nm-scale oscillations. We know this to be a realistic vibration length scale, and furthermore, we know that the vibrations are dominated by low-frequency acoustic pick-up transferred through the helium dewar (Sec. 3.1.1). If this hypothesis is correct, the

B-field noise can be improved in a number of ways. First of all, magnets can be attached directly to the 4K chamber, which would cause them to vibrate common-mode with the trap. Alternatively, a Helmholtz-like configuration (Sec. 3.4.3) could be employed once again and combined with more precise coil positioning.

## 5.5. Rabi oscillations

### 5.5.1. Optical qubit rotations

Rabi oscillations can be driven by preparing the qubit in  $|1\rangle$  and applying a pulse of  $\lambda = 729\text{ nm}$  light on resonance with the optical qubit transition  $|S_{1/2}, m_J = -1/2\rangle \rightarrow |D_{5/2}, m_J = -1/2\rangle$ . The resulting Rabi oscillations are shown in Fig. 5.21. Fitting the



**Figure 5.21:** Rabi oscillations on the  $|1\rangle = |S_{1/2}, m_J = -1/2\rangle \rightarrow |0\rangle = |D_{5/2}, m_J = -1/2\rangle$  transition. (left) Rabi oscillations following Doppler cooling of all motional modes (right) Rabi oscillations following Doppler cooling of all motional modes and EIT cooling the axial mode.

oscillation frequency allows us to deduce the amplitude of the applied electric field  $E_0$  (Eq. (2.9)). Assuming the ion is at the centre of a Gaussian beam with waist size  $w_y \times w_x$ , the beam power can be calculated as:

$$P_{\text{ion}} = \frac{E_0^2 \pi w_x w_y}{2\epsilon_0 c},$$

which can be then compared to the input power in the fibre  $P_{\text{input}}$ . Alternatively, we can estimate  $P_{\text{ion}}$  from the measured loop transmission and the simulated grating efficiency. These measurements were performed the trap #4, where we measured a pi-time  $t_\pi = 2.4\text{ }\mu\text{s}$



for the input power of  $P_{\text{input}} = 1.5 \text{ mW}$ . This corresponds to:

$$\begin{aligned} P_{\text{ion}}/P_{\text{input}} &= 0.14 && \text{Rabi frequency fit,} \\ P_{\text{ion}}/P_{\text{input}} &= 0.23 && \text{Loss measurement + simulation.} \end{aligned}$$

The likely explanation for this discrepancy is an offset between the micromotion null and the beam centre. The whole difference could be accounted for by a  $1.8 \mu\text{m}$  offset along  $y$  or by  $3.1 \mu\text{m}$  offset along  $x$ . This is consistent with the specified and measured misalignment of the waveguide layer of  $< 2 \mu\text{m}$ . However, measurements in other zones did not produce higher Rabi frequencies, ruling out a strong misalignment along  $y$  only. The second possible contributing factor is the variability between channel-to-channel fibre coupling efficiencies. Specifically, different fibres experience a different loss on the vacuum fibre feedthrough, and a  $\pm 10\%$  loss variability per connector is not atypical. Finally, the expression for the Rabi frequency in Eq. (2.9) is only valid for a monochromatic plane wave, and we discovered it fails to capture the relative Rabi frequencies of different carriers observed in our experiment (Sec. 9.3.5). C. Zhang performed a heroic calculation of Rabi frequencies which takes into account the full beam profile (Sec. 4.30), and further details can be found in his thesis [Zha21].

### 5.5.2. Rabi frequency noise


The decay profile of Doppler-cooled oscillations is dominated by the axial mode temperature. By EIT-cooling this mode from  $\bar{n} \approx 8$  to  $\bar{n} < 0.5$ , Rabi-oscillations increase in contrast and speed up by  $\approx 2\%$ . The Gaussian fits (Fig. 5.21) imply the standard deviation of the the Rabi frequency of :

$$\begin{aligned} \sigma_{\Omega_{\bar{n}}}/\langle\Omega_{\bar{n}}\rangle &= 2.5(1) \times 10^{-2} && \text{Doppler,} \\ \sigma_{\Omega_{\bar{n}}}/\langle\Omega_{\bar{n}}\rangle &= 1.3(1) \times 10^{-2} && \text{Doppler + EIT.} \end{aligned}$$


The residual Rabi frequency fluctuations are not inconsistent with the measured temperatures of the Doppler-cooled radial modes ( $\bar{n} \approx 6.7$  for the lower radial, and  $\bar{n} \approx 2.3$  for the upper radial) (Sec. 5.6.2), although we did not keep track of their respective Lamb-Dicke parameters systematically. The relation between the mode temperature and the residual Rabi frequency error is plotted in Fig. 5.22. Given the observed Rabi frequency fluctuations, we expect the Rabi oscillations to reach a depth of  $P(1) = 4(1) \times 10^{-4}$ , but experimentally, they saturate at a minimum value of  $P(1) = 1.0(1) \times 10^{-3}$ . This is more than can be accounted for by the state preparation error (Sec. 5.2) and the readout error (Sec. 5.3), and additional experiments are necessary to understand the discrepancy.

Further insight about the origin of Rabi frequency noise can be gained with an experiment shown in Fig. 5.23. A detuned  $\lambda = 729 \text{ nm}$  is inserted into the wait time of a Ramsey sequence, which creates a power-dependent Stark shift. By using a spin-echo sequence, with

## 5. SINGLE-ION OPERATIONS WITH INTEGRATED WAVEGUIDES



**Figure 5.22:** Rabi frequency error caused by thermally occupied modes of a single-ion calculated from Eq. (2.38). The error is expressed as the standard deviation in the Rabi frequency divided by the mean Rabi frequency, assuming a Lamb-Dicke parameter given by Eq. (2.12). The angle  $\theta$  is known for the axial mode, but not for the radial modes. As a result, the radial mode curves only serve as an example and use  $\theta = 45$  deg.




**Figure 5.23:** Power stability measurement scheme. A far-detuned  $\lambda = 729$  nm pulse is inserted into one arm of a spin-echo Ramsey experiment. The final  $\pi/2$ -pulse measures the phase induced by the laser Stark shift, indirectly measuring the Rabi frequency.

a laser pulse inserted into only one branch, we can cancel out the static qubit detuning errors, simplifying the experimental procedure. Using the analysis of Sec. 5.4.2 we find:

$$P(1) = \frac{1}{2}(1 - \cos(\Delta\phi)) \quad \text{where} \quad \Delta\phi = \int_0^\tau \frac{\Omega^2(t)}{4\delta(t)} dt,$$

where  $\Omega(t)$  is the Rabi frequency and  $\delta(t)$  is the detuning of the spectator transition. By choosing a large enough value of  $\delta$ , the fluctuations in  $\Delta\phi$  are dominated by fluctuations in  $\Omega(t)$ , and detuning noise can be neglected.

We perform the experiment in two ways. In the first, the Ramsey experiment is performed on the Zeeman qubit, and the  $\lambda = 729$  nm laser is near-resonant with an optical qubit transition. In this scenario,  $\Omega$  is the Rabi frequency of the nearest quadrupole transitions, and the dipole Stark shifts do not create an observable phase shift. In the second way, the Ramsey experiment is performed on the optical qubit, and the  $\lambda = 729$  nm is far-detuned any quadrupole transition ( $\delta \approx 2\pi \times 300$  MHz). In this case, the Stark shift is dominated by couplings to the dipole transitions. The comparison between the two sequences is shown in Fig. 5.24. The Dipole Stark shift exhibits much smaller fluctuations than the quadrupole



**Figure 5.24:** Power stability measurement results (left) The quadrupole Stark shift, measured on a Zeeman qubit (right) The dipole Stark shift, measured in the optical qubit.

Stark shift. The Gaussian fits in Fig. 5.24 return:

$$\begin{aligned} \sigma_{\Omega}/\langle\Omega\rangle &= 1.3(1) \times 10^{-2} && \text{Quadrupole Rabi frequency, near-detuned} \\ \sigma_{\Omega}/\langle\Omega\rangle &= 1.1(1) \times 10^{-3} && \text{Dipole Rabi frequency, far-detuned.} \end{aligned}$$


The quadrupole Rabi frequency fluctuations are consistent with those obtained from Rabi oscillations in Sec.5.21. On the other hand, the dipole Rabi frequency is independent of the ion's temperature and thus is immune from the main source of the quadrupole Rabi frequency fluctuations. The residual noise is then directly related to the drive intensity drifts, which can originate from the ion motion within the beam, or from drifts of the input power or polarization.

To understand the remaining noise on the dipole Rabi frequency, we apply a  $\lambda = 729$  nm pulse to both arms of the spin-echo sequence. In this case, the contrast is indistinguishable from that of the Ramsey experiment alone (Fig. 5.25). This implies that the vast majority of the observed noise is slow, as would be expected e.g. for polarization drifts. Since the contrast decay is dominated by the laser frequency noise, we cannot measure the remaining Stark shift noise with high precision. The decay curves are identical within the error bars of  $\pm 2\%$ , which suggests the residual fast noise at a level of  $\sigma_{\Omega}/\langle\Omega\rangle \lesssim 10^{-5}$ .

### 5.5.3. Zeeman qubit rotations


Zeeman qubit Rabi oscillations are performed by feeding a current oscillating at the  $|S_{1/2}, m_J = -1/2\rangle \leftrightarrow |S_{1/2}, m_J = 1/2\rangle$  splitting frequency of  $\omega = 2\pi \times 16.5$  MHz through a track on the trap carrier PCB. We typically operate at a pi-time of  $t_{\pi} \approx 6.4$   $\mu$ s, limited

## 5. SINGLE-ION OPERATIONS WITH INTEGRATED WAVEGUIDES



**Figure 5.25:** Contrast obtained in an optical qubit Ramsey sequence with spin echo and no far-detuned light (red) is statistically indistinguishable from that obtained with inserting a far-detuned pulse into both arms of the power stability experiment in Fig. 5.23. This indicates low contribution of high-bandwidth noise to the dipole Rabi frequency error.

by the available amplifier. Precise pi-time calibration can be made by applying the pulse multiple times in a row [Hom06]. Initial experiments with up to 11 pi-pulses revealed a strong dependence between the average Rabi frequency and the duty cycle of the experiment, as shown in Fig. 5.26 (red points). We speculated that this error was caused by poor




**Figure 5.26:** Zeeman qubit pi-time calibration from 11 pulses. Initially, the system exhibits strong transients, with a strong correlation between pulse area and the duty cycle (red points). After adding an attenuator to the RF chain, the duty-cycle dependence becomes negligible (blue points).

impedance matching between the signal source and the bandpass filter. Reflections from the filter would return to the amplifier and circulate in the system over  $\approx 10 \mu\text{s}$ . To test this hypothesis, we replaced the amplifier with a stronger one, followed by an attenuator to keep the overall gain constant. We speculated that the attenuator would lower the overall Q-factor of the circuit, increasing the dissipation rate of the reflected power. Indeed,

following that fix, the duty cycle issues were not visible down to the minimum gap length of  $1.4\,\mu\text{s}$  (Fig. 5.26, blue points).

Since the Lamb-Dicke parameter of the drive is negligible, the Rabi frequency does not exhibit thermal shot-to-shot fluctuations. Consequently, we expected the Rabi oscillations to exhibit much longer coherence. As shown in Fig. 5.27, that is indeed the case, and a few hundred oscillation cycles can be applied before the contrast begins to decay. While



**Figure 5.27:** Contrast of Zeeman qubit Rabi oscillations vs number of Rabi oscillations. The solid line shows a rough Gaussian fit to the decay.


the Gaussian fit is not excellent, it suggests an approximate strength of shot-to-shot Rabi frequency fluctuations of:

$$\sigma_{\Omega}/\langle\Omega\rangle \approx 1.3 \times 10^{-4}$$

Further examination suggests that the dominant source of amplitude noise is slow and related to the amplifier duty cycle. This can be seen by comparing the Rabi oscillation curve around  $t = 1000\,\mu\text{s}$  ( $t/t_{\pi} = 156$ ) with the one for  $t = 5000\,\mu\text{s}$  ( $t/t_{\pi} = 781$ ), shown in Fig. 5.28. In both cases, the data points were acquired from left to right. While the curve on the right fits with unity contrast, the three left-most data points are clearly offset from the fit curve. The most likely hypothesis is that at the beginning of the experiment (when the duty cycle changes) the Rabi frequency drifts over the  $\approx 1\,\text{s}$ , which is the time necessary to acquire the first few data points. Afterwards, the amplifier has reached the steady state and the subsequent points all have the same Rabi frequency. If this hypothesis is correct, Zeeman qubit rotation error could be decreased by duty cycle stabilisation.

## 5.6. Cooling and heating

This section discusses our ability to keep the ion close to the motional ground state. Sec. 5.6.1 describes the methods used to diagnose the temperature of different modes of motion. The



**Figure 5.28:** Comparison of Zeeman qubit Rabi oscillations at different times pulse times.

next two sections describe the calibration and performance of the two cooling methods we most commonly use in single-ion experiments: Doppler cooling and EIT cooling. Following ground-state cooling, the ion heats up at a constant rate due to ambient electric-field noise. We measure this heating rate by cooling the ion and measuring its temperature following a variable-length delay. Those measurements are presented and interpreted in Sec. 5.6.4.

### 5.6.1. Ion thermometry

Throughout this thesis, we used three different methods of estimating mode temperatures. In this section, we compare the different methods and their merits.

#### Sideband fitting method

Consider an ion prepared in an internal state  $|1\rangle$  and a thermal state with mean phonon number  $\bar{n}$ . The ion is driven with a pulse resonant with the BSB transition. The resulting sideband Rabi oscillations take a form:

$$P_B(0) = \sum_{n=0}^{\infty} P_{\text{th}}(n) \sin^2 \left( \eta \sqrt{n+1} \Omega t / 2 \right) \quad (5.6)$$

where  $P_{\text{th}}(n)$  is given by Eq. (2.37),  $\Omega$  is the carrier Rabi frequency and  $\eta$  is the Lamb-Dicke parameter (Eq. (2.12)) for the mode in question. One of the most straight-forward methods of estimating  $\bar{n}$  is to drive the sideband oscillations and fit them with Eq. (5.6). Most of the time, the essential information is contained in the time of the first maximum of  $P(0)$ , which is  $t_{\pi,bsb} \approx \pi / \Omega \eta \sqrt{\bar{n} + 1}$ .

In order to extract the temperature, it is necessary to know  $\eta$ , and hence the angle between the laser and the mode eigenvector. This necessitates additional measurements for the radial modes, whose orientation is not set by the trap geometry. Another downside of

the method is that the sideband oscillations rarely fit the form of Eq. (5.6) exactly, and additional decoherence parameters have to be introduced to explain the data [Mee96]. We routinely use the sideband fitting method to extract the thermal occupation of the axial mode at low to moderate temperatures ( $\bar{n} \approx 0 - 20$ ).

### Sideband ratio method

Suppose that in addition to measuring the BSB Rabi oscillations, we measure the RSB oscillations in a separate experiment. The RSB trace is given by

$$P_R(0) = \sum_{n=1}^{\infty} P_{\text{th}}(n) \sin^2(\eta\sqrt{n}\Omega t/2). \quad (5.7)$$

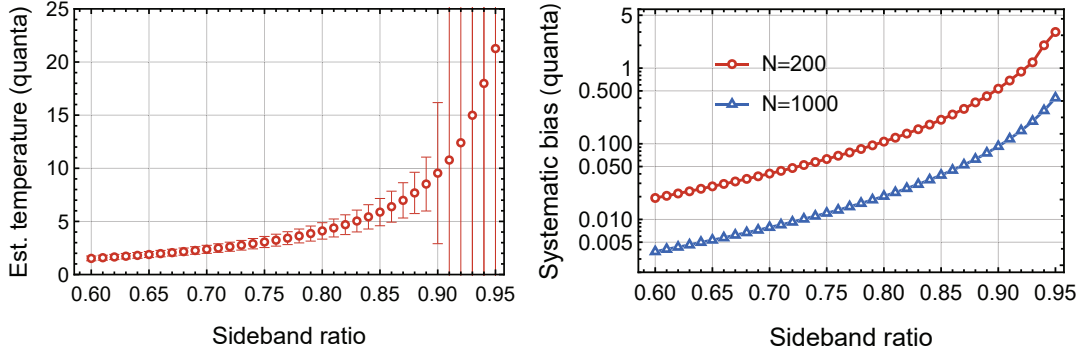
The ratio of the red and blue sideband excitation probabilities can be directly evaluated from Eq. (5.6) and Eq. (5.7) to read [Hom06]:

$$R = \frac{P_R(0)}{P_B(0)} = \frac{\bar{n}}{\bar{n} + 1}.$$

Remarkably,  $R$  is independent of both the excitation time and the Lamb-Dicke parameter! This allows us to estimate the temperature directly from the two experiments, without fits or additional input parameters. Furthermore, the ratio test provides immunity from additional common-mode decoherence of sideband oscillations. The temperature can be estimated from

$$\bar{n} = \frac{R}{1 - R}.$$

The biggest downside of the sideband ratio method is that it can only be robustly applied at low thermal occupations ( $n \lesssim 5$  or so), and becomes a statistics nightmare as  $R \rightarrow 1$ . The problem is that while  $R$  is generally normally distributed (aside from  $R \approx 0$ ), the estimate of  $\bar{n}$  is not. As  $R \rightarrow 1$ , the standard deviation in the estimate of  $\bar{n}$  increases dramatically, as shown in Fig. 5.29 (left). This process introduces a lot of fluctuations in the experimental results unless a large amount of data is averaged together. The second problem is that  $R/(1 - R)$  is a biased estimator of the  $\bar{n}$ , as shown in Fig. 5.29 (right). In other words, the measured average of  $R/(1 - R)$  systematically overestimates the true value of  $\bar{n}$  by an amount that depends on the number of averaged data points. The final issue of the sideband ratio method is that, at higher temperatures, it may become biased due to the different dephasing rates of different Fock states [Kin98]. That being said, we routinely use the sideband ratio method to estimate the temperatures of all three motional modes in the range of  $\bar{n} = 0 - 5$ . We also employ it for heating rate measurements of the axial mode, assuming it is first ground-state cooled.



**Figure 5.29:** Simulated issues with the sideband ratio method. (left) As  $R \rightarrow 1$ , the standard deviation in the estimate of  $\bar{n}$  explodes. This is illustrated for  $N = 200$  experiments per point (right)  $R/(1 - R)$  is a biased estimator for  $\bar{n}$ . The amount of bias depends on both on  $R$  as well as on the number  $N$  of experimental repetitions, with more repetitions resulting in lower bias.

### Carrier slow-down method

The final method estimates the ion temperature from the Rabi frequency of carrier oscillations (Sec. 2.5.1). While this correction is very small when the motional modes are close to the ground state, it becomes significant at larger temperatures, allowing for reliable thermometry even at  $\bar{n} > 100$  [Noe19]. For a measured Rabi frequency of  $\langle \Omega_{\bar{n}} \rangle$ , and given a Rabi frequency  $\Omega_0$  measured at close to motional ground state, the temperature can be estimated from (Eq. (2.38)):

$$\langle \Omega_{\bar{n}} \rangle / \Omega_0 = \prod_i e^{-\eta_i^2 \bar{n}_i}, \quad (5.8)$$


where the product acts over all motional modes. The carrier Rabi frequency reflects the temperature of all the motional modes and does not allow for the determination of the individual values of  $\bar{n}_i$ . However, it is common that the carrier slow-down is dominated by the lowest-frequency mode. This is because that mode has the highest Doppler temperature (Sec. 2.5.2), the largest heating rate and the largest Lamb-Dicke parameter. We determined through the sideband ratio method that the radial modes indeed exhibit insignificant heating compared to the axial mode (Sec. 5.6.4). Therefore, when analysing the carrier slow-down data, we assume the right-hand side of Eq. (5.8) to be equal to  $\exp(-\eta^2 \bar{n})$  for the axial mode. As a result, we expect the result to slightly overestimate the true  $\bar{n}$ .

#### 5.6.2. Doppler cooling

We optimise Doppler cooling by adjusting the amplitude and detuning of the  $\lambda = 397$  nm  $\pi$  laser, as well as the overall cooling time. Since the Doppler limit (Sec. 2.5.2) is only achieved in the limit of low power and long cooling time, we adjust the  $\lambda = 397$  nm  $\pi$  power below the saturation intensity and measure the necessary cooling duration to reach the steady-state. We monitor the final temperature by measuring the sideband ratio of the



upper radial mode at  $\omega = 2\pi \times 5.2$  MHz. As the highest frequency mode, it has the lowest Doppler temperature, which makes the ratio test reliable. Fig 5.30 summarises the cooling performance. The cooling cycle takes  $\approx 1 - 2$  ms to reach a steady state. The optimal



**Figure 5.30:** Sideband ratio of the upper radial mode at  $\omega = 2\pi \times 5.2$  MHz vs Doppler cooling parameters

cooling is achieved at  $\delta \approx -2\pi \times 8$  MHz detuning from the detection resonance, in agreement with the theoretical value of  $\delta = -\Gamma/2$  for a two-level system. The sideband ratio method implies the final mode temperature of  $\bar{n} \approx 2.3$ , matching the theoretical Doppler limit for a mode angle  $\theta_2 \approx 60$  deg.

Following the optimisation, we measure the temperatures of the other motional modes. We find  $\bar{n} \approx 8$  for the axial mode at  $\omega = 2\pi \times 1.3$  MHz, again in agreement with the Doppler limit. However, for the lower radial mode at  $\omega = 2\pi \times 3.6$  MHz we find  $\bar{n} \approx 6.5$ , which suggests a large angle of  $\theta_1 \approx 70$  deg between the mode vector and the cooling laser.


Independent measurements of the radial mode Lamb-Dicke parameters imply that the lower and upper radial modes are tilted by 25 deg and 65 deg respectively from the trap normal. The fact that these numbers do not sum to 90 deg suggests an  $\pm 5$  deg measurement error. The measured tilts correspond to angles of  $\theta_1 = 72(3)$  deg and  $\theta_2 = 52(2)$  deg with respect to the cooling beam (which is parallel to the trap and at 45 degrees to the trap axis), in qualitative agreement with those obtained above. This suggests that the radial temperature can be decreased by increasing the mode tilt from the vertical. We briefly attempted to do that, but haven't found a better operating position yet. Nonetheless, the measured temperatures are in respectable agreement with a simple two-level system model.

### 5.6.3. EIT cooling

In our experiment, the  $\lambda = 397$  nm  $\pi$  and  $\sigma$  lasers are both oriented at 45 degrees to the trap axis, with their wavevector difference pointing along the trap axis. As a result, EIT light only cools to the axial modes, leaving the radial modes unaffected.

## 5. SINGLE-ION OPERATIONS WITH INTEGRATED WAVEGUIDES

EIT cooling can be optimised as follows. We set the  $397\sigma$  detuning  $\delta_\sigma \approx 2\pi \times 45$  MHz. The  $397\sigma$  laser is optimised for high efficiency at this detuning, while the  $397\pi$  drive still has sufficient bandwidth to act as a probe beam. We then apply a BSB pulse with duration  $2t_{\pi,bsb} = 2\pi/\Omega\eta$ , which would drive an ion prepared in  $|1\rangle$  and  $\bar{n} = 0$  back to  $|1\rangle$ . We scan the amplitude of the  $397\sigma$  laser together with the frequency offset of the  $397\pi$  probe beam from the  $397\sigma$  drive. For the correct set of settings, we observe a clear increase in  $P_B(1)$ , as shown in Fig. 5.31 (left). Fig. 5.31 (right) shows a time trace of the BSB oscillations at the



**Figure 5.31:** EIT cooling calibration and performance (left)  $P(1)$  at  $2t_{\pi,bsb}$  vs  $397\sigma$  amplitude and  $397\pi$  detuning. Well-calibrated cooling creates a higher value of  $P(1)$  (red regions). (right) BSB oscillation on an EIT-cooled mode. The theoretical curve corresponds to unitary evolution for  $\bar{n} = 0.3$ .

for optimal settings, together with a fit from Eq. (5.6) with  $\bar{n} = 0.3$ . After  $2t_{\pi,bsb}$  we begin to notice significant discrepancies between the data and the fit, probably due to motional dephasing and heating. The total duration of the EIT cooling sequence is 150 μs.

The EIT offers a robust way to cool the axial mode to  $\bar{n} < 1$ , and the sequence does not require frequent calibrations. However, the exact value of the final temperature is still a bit of a mystery. In the setup 1.0, we have achieved temperatures as low as  $\bar{n} = 0.1$ , but  $\bar{n} = 0.3$  was more common. In the setup 2.0, we sometimes measure  $\bar{n} = 0.3$ , but more commonly  $\bar{n} \approx 0.6$ , and the final temperature is not clearly correlated with the heating rate.

### 5.6.4. Heating rates


Following ground-state cooling, the three modes of motion begin to heat up at different rates. A mode at a frequency  $\omega$  heats due to electric-field noise at that frequency. Formally, the heating rate  $\dot{n}$  from the ground state is given by

$$\dot{n}(\omega) = \frac{e^2}{4m\hbar\omega} S_E(\omega), \quad (5.9)$$

where  $m$  is the ion mass and  $S_E(\omega) = 2 \int_{-\infty}^{+\infty} dt \langle E(t)E(0) \rangle e^{i\omega t}$  is the electric-field noise density caused by electric field component  $E(t)$  along the mode in question [Bro14]. In general, higher-frequency modes heat up at a slower rate than the low-frequency ones. This is both because of the  $1/\omega$  term in Eq. (5.9), and because most noise sources produce less electric field noise at higher frequencies.

Using as a reference other cryogenic SETs with 50  $\mu\text{m}$  ion-electrode distance, we can expect an electric field noise density in the range of  $S_E(\omega) = 5 \times 10^{-14} - 1 \times 10^{-12} \text{ V}^2/\text{m}^2/\text{Hz}$ , corresponding to a heating rate of  $\dot{n} \approx 10 - 100$  quanta/s at  $\omega = 2\pi \times 1.0 \text{ MHz}$  [Chi14a; Bru15; Sed18b]. However, the measured heating rates are one to two orders of magnitude larger. The measurements presented in this section attempt to understand the origin of this excess noise.

### Frequency dependence




**Figure 5.32:** Heating rate vs frequency. The heating rates are broadly consistent with a  $1/f^2$  trendline (dashed line, not a fit), corresponding to the  $1/f$  dependence of the electric field noise spectral density. Except for the two measurements on the radial modes (highlighted with arrows), all other were performed on the axial mode. The axial mode heating rate plateau above  $f \approx 1.5 \text{ MHz}$  deviates significantly from the trendline. The increase of the heating rate below 1.5 MHz was observed in trap #5 as well, but not studied systematically, while for the traps #2 and #4 the axial frequency range was limited by the available electrode voltages and the strong static curvature of the trap. The lower frequency radial mode is closer to the trap surface normal than the higher frequency radial mode.

Fig. 5.32 shows the measured heating rates for various waveguide-integrated traps, different modes and different frequencies. For the axial mode, we heating rates decrease with increasing mode frequencies up to  $\omega \approx 1.5 \text{ MHz}$ , at which point they appear frequency-insensitive. The plateau occurs at a heating rate of  $1000 - 2000$  quanta per second, corresponding to electric field noise in excess of  $2 \times 10^{-11} \text{ V}^2/\text{m}^2/\text{Hz}$ . This is over  $100\times$  higher than what the best results reported for identical ion-electrode distance [Chi14a] and similar to the heating rates observed in room-temperature traps [An19].

Further increasing the axial mode frequency increase does not lower the heating rate, which is particularly surprising given that the radial modes exhibit a significantly lower heating rate. The latter are consistent with a general trend of  $S_E(\omega) \propto 1/\omega$  (although we did not attempt to measure the exponent precisely)

### Temperature dependence

The trap can be locally heated through a resistor on the trap carrier PCB, which also serves as a temperature sensor. However, the resistor calibration is only specified down to 77 K and thus the trap temperature cannot be directly measured for low heater powers. In this intermediate regime, we estimate it as follows. With the heater turned off, we assume that the trap is in thermal equilibrium with the cryogenic chamber, whose temperature of 7 K is measured with a second calibrated sensor. We then turn on the trap heater and increase its power until the resistance is within the calibrated range. In our case, we measure the trap temperature of 150 K for 400 mW dissipated on the heater. Assuming the power is dissipated predominantly by thermal conduction across the trap interposer to a constant-temperature reservoir, the trap temperature increase should be proportional to the heater power. This way, we can estimate the trap temperature of  $\approx 20$  K for a heater power of 40 mW. Alternatively, we can use example literature values of the temperature dependence of the heater resistance and apply them to our measurements [Ram14]. These suggest the trap temperature at 40 mW is closer to 30 K.



**Figure 5.33:** Axial heating rate vs trap chip temperature. The temperature is controlled by setting a fixed trap heater power (bottom). The chip temperature is then estimated from the heater resistance measurements (top), see main text.

Using this approximate calibration, we measure the temperature dependence of axial heating rates shown in Fig. 5.33. We find the heating rate increases by an order of magnitude between 7 K and 150 K. Despite the approximate low-temperature calibration, we find a clear noise increase by a factor of 2-3 for temperature increases as modest as 10 – 20 K.

This is qualitatively similar to the behaviour reported in [Sed18b] for ion trap without surface treatment, albeit at a much higher overall electric-field noise density.

### Possible causes

The origin of anomalously high heating rates is currently unknown, but the reported measurements can guide us towards certain hypotheses. The observation of large increases in the electric-field noise density even for small temperature changes suggests the trap is the origin of the noise. The magnitude of that electric-field noise change exceeds the values reported in the literature, suggesting a novel mechanism is at play.

Of course, without cooling the trap further, it is challenging to definitely assert whether the large heating rate at 7 K is also due to this thermally-activated mechanism, or whether there is another cause. To that end, we performed an exhaustive search of possible technical noise sources. The heating rates are largely independent of the applied RF power or radial stray fields, suggesting that RF noise is not the culprit. Adding or removing additional ground connections or ground loops did not influence the electric field noise in a meaningful way either. We also recorded the heating rates with unnecessary laboratory devices switched off, disconnected or terminated. This included measuring the electric field noise with the RF source as the only cable connected to the experiment (the DC lines can be unplugged as the electrode voltages take many hours to discharge through the filterboard capacitors). We also performed the measurements with the cryostat temporarily switched off to rule out any noise induced by the compressor electronics. In all the cases, the heating rates remained unchanged. This suggests, but does not prove, that technical noise is a small contribution to the overall noise budget<sup>5</sup>.

Further (but still weak) evidence against the technical noise hypothesis is the lack of strong noise polarization [An19]. Although we have not simulated it for our trap, one would expect the technical noise to be polarized primarily along the trap normal, especially if the noise is common-mode to all electrodes, as could be expected e.g. for electromagnetic pickup [Sch15]. In that case, one would expect to observe reduced electric-field noise density for the modes oriented closer to the trap parallel (axial and upper radial), and increased electric-field noise on the mode oriented closer to the trap normal (i.e. the lower radial mode, which is at  $\approx 25^\circ$  to the trap normal). However, while all the modes follow the same broad trend of  $S_E(\omega) \propto 1/\omega$ , the main outliers above the trendline are high-frequency axial modes. Thus, it appears the noise is not predominantly along the trap normal, slightly strengthening the trap-embedded fluctuator hypothesis.

While I do not know what causes the overall heating rate, it is worth pointing out a few unusual trap features as potential suspects. Firstly, the ion has a direct line of sight to a large amount of dielectric material, as well as to the silicon substrate (at the grating

---

<sup>5</sup> If future experiments eliminate the trap-related noise and become limited by technical noise once again, I would suggest investigating the effect of shielding the DC cables, which are currently susceptible to electrical pick-up inside the vacuum system.

## 5. SINGLE-ION OPERATIONS WITH INTEGRATED WAVEGUIDES

couplers). At the time of writing this thesis, it was pointed out that an ion will necessarily heat up in the presence of nearby dielectrics due to coupling to image charges [Tel21]. I have not simulated the influence of this effect on our system, although I expect it to be significantly attenuated at cryogenic temperatures. It is worth pointing out that the heating rates measured in the trap #5 are consistent with those measured in the trap #2, which had no integrated waveguides (Fig. 5.32). This suggests that the presence of integrated optics is not the cause of enhanced electric-field noise. The possible influence of dielectrics and silicon on the heating rates could be constrained by comparing the heating rates close to and far away from the grating couplers. However, the ion has significant exposure to substrate dielectrics everywhere along the trap axis. It would be therefore interesting to measure whether the heating rates improve in future fabrication runs with additional substrate shielding.

Secondly, the ground plane separating the trap electrodes from the silicon substrate is relatively resistive, with  $\sim 1\ \Omega$  resistance from one corner of ground plane to the other, and  $> 3\ \Omega$  resistance per via to the top metal layer. It would be interesting to measure whether the heating rates improve in future fabrication runs with increased ground-plane conductivity.

## 6. Quantum contextuality

Experiments presented in the next chapter describe experimental observations of *quantum contextuality* through violations of *non-contextuality (NC) inequalities*. This chapter describes the theory underlying those experiments.

The first stumbling block on this experimental journey was understanding what quantum contextuality actually means, and how the theoretical concepts can be reconciled with experimental realities. There was a lack of relevant review papers or lecture notes (although an extensive review was published after this section was written [Bud21]), and the published literature was often contradictory and opinionated. Experimental papers, on the other hand, tended not to delve deep into the issues of interpretation. Through numerous conversations, I found out that, in this field, many experimental physicists feel confused by the theoretical advice, and resort to performing the experiments without asking too many questions.

### Warning

This chapter is written for an experimental physicist who aims to understand quantum contextuality. Unlike the other chapters, it aims to describe the main results in the style of lecture notes, including theoretical prerequisites and complete derivations. Most of the chapter does not present original work but gathers calculations from various sources to create a single narrative.

A reader uninterested in the theory and philosophy of contextuality tests will find the chapter tedious, and is discouraged from reading it all. Instead, the following sections are sufficient to understand the experiments in Chapter 7:

1. KCBS inequality (Sec. 6.3.1) and its quantum violation (Sec. 6.3.2).
2. Generalised KCBS inequality (Sec. 6.3.4).
3. Yu-Oh inequality (Sec. 6.4.1).
4. Experimental loopholes (Sec. 6.5).

I believe that the best way to introduce NC inequalities is a generalisation of the Bell inequalities. The reason for this is two-fold. First of all, Bell inequalities are more familiar to most physicists, and I hope that by presenting quantum contextuality in the same framework, I will manage to de-mistify many traditionally obscure concepts. Secondly, a lot of quantum contextuality theory takes the concepts that naturally apply to non-local systems and adapts them (with varying degrees of success and clarity) to local systems.

I hope that this angle of presentation will shed some clarity and motivation onto the experiments presented in the next chapter.

This review is structured as follows. Sec. 6.1 introduces the quantum information formalism necessary to derive the Bell and NC inequalities. Specifically, it formalises the concepts of quantum entanglement and quantum measurements. In Sec. 6.2 I derive two versions of Bell’s theorem, which show the non-classical nature of entangled states. Then I introduce the idea of quantum contextuality, which aims at extending Bell’s insights onto local systems. This is done through a derivation of the KCBS inequality, which will be experimentally verified in Chapter 7. Sec. 6.3.4 connects the KCBS inequality and its extension of the chained Bell’s inequalities, strengthening the connection between local and non-local contextuality. Sec. 6.4 introduces the concept of state-independent contextuality and the Yu-Oh inequality which will be experimentally verified in Chapter 7. Sec. 6.5 describes how the theoretical concepts can be applied in experiments. Specifically, we review a long-lasting debate on the so-called “compatibility loophole”, and the approaches used to overcome it. The final section presents my personal contribution and comments. It attempts to give an intuition about the meaning of contextual and non-contextual models (Sec. 6.6.1) and constructs explicit examples of classical noise sources which can masquerade as non-classical corrections in contextuality test (Sec. 6.6.3). It then discusses what can be learned from contextuality experiments (Sec. 6.6.4), and criticises the idea of using contextuality for randomness generation (Sec. 6.6.5).

## 6.1. Quantum measurements and entanglement

Both nonlocality and contextuality tests rely on correlations of quantum measurements. In the former case, these measurements are performed on separate particles of a bipartite entangled system, while in the latter, the measurements are performed sequentially on a single system. In this section, we briefly review the formalism of quantum entanglement and quantum measurements.

### 6.1.1. Composite and entangled states

Consider a quantum system  $A$  with a set of orthonormal basis states  $\{|i\rangle_A\}$ . Any pure state  $|\psi\rangle_A$  of  $A$  can be written as

$$|\psi\rangle_A = \sum_i c_i^A |i\rangle_A, \quad (6.1)$$

where  $c_i^A$  is a unit vector. Likewise, any pure state  $|\phi\rangle_B$  of a different quantum system  $B$  with basis states  $|i\rangle_B$  can be written as

$$|\phi\rangle_B = \sum_i c_i^B |i\rangle_B, \quad (6.2)$$



where  $c_i^B$  is a unit vector. Let us now consider  $A$  and  $B$  as parts of one system. In this composite system, the basis states are  $|i\rangle_A \otimes |j\rangle_B$ . Consequently, any pure state can be written as:

$$|\chi\rangle_{AB} = \sum_{i,j} c_{i,j} |i\rangle_A \otimes |j\rangle_B. \quad (6.3)$$

### Entanglement

The composite state in Eq. (6.3) is *separable* if and only if it can be written as a tensor product of pure states of  $A$  and  $B$ , that is:

$$|\chi\rangle_{AB} = |\psi\rangle_A \otimes |\phi\rangle_B. \quad (6.4)$$

Plugging Equations (6.1) and (6.2) into Eq. (6.4) and comparing with Eq. (6.3) implies that an equivalent condition for separability is the existence of coefficients  $c_i^A$  and  $c_i^B$  which satisfy

$$c_i^A c_j^B = c_{i,j}$$

for all  $i$  and  $j$ . All bi-partite states which do not satisfy this condition - that is, are not separable - are said to be *entangled* [Nie02]. The intuition is that we can think of separable states by considering their parts separately, while entangled states only make sense when we look at them as a whole.

We can return to a local sub-system description by introducing a *reduced density matrix*  $\rho_A$ . Formally,  $\rho_A$  is obtained by the partial trace operator acting on the basis states of  $B$ :

$$\rho_A = \text{Tr}_B \rho_{AB} = \sum_j {}_B \langle j | \rho_{AB} | j \rangle_B.$$

The state  $|\chi\rangle_{AB}$  is *maximally entangled* if and only if

$$\rho_A = \frac{1}{N} \times \mathbb{1},$$

where  $N$  is the dimension of  $\rho_A$ . It can be shown that the same condition applies to  $\rho_B = \text{Tr}_A \rho_{AB}$ .

In order to gain an intuitive understanding of this definition, let us introduce the concept of *von Neumann entropy* of a density matrix:

$$S(\rho) = -\text{Tr}(\rho \ln(\rho)).$$

$S(\rho) = 0$  for pure states, increases of mixed states, and reaches a maximum value of  $\ln(N)$  for a maximally mixed state  $\rho = \frac{1}{N} \times \mathbb{1}$  [Ben06]. Intuitively, measurements of a pure state generate the least uncertainty, since one can pick a measurement basis where measurement

## 6. QUANTUM CONTEXTUALITY

outcomes are deterministic. On the other hand, measurements of a mixed state always generate random outcomes.

Armed with this understanding, we see that maximally entangled states satisfy  $S(\rho_{AB}) = 0$  and  $S(\rho_A) = S(\rho_B) = \ln(N)$ . In other words, the system as a whole ( $\rho_{AB}$ ) is minimally random, but its sub-systems ( $\rho_A$  and  $\rho_B$ ), taken on their own, look maximally random. So for maximally entangled states, all of the information is shared, and none is accessible locally.

### Bell states

Let us now limit the discussion to bipartite qubit states. In order to find maximally entangled states of two qubits, we solve a general two-qubit pure state in Eq. (6.3) to have a fully mixed reduced density matrix  $\rho_A = \frac{1}{2} \times \mathbb{1}$ . From this condition, it is easy to show that all two-qubit maximally entangled states have one of the following forms

$$\begin{aligned} |\Phi(\phi)\rangle &= \frac{1}{\sqrt{2}} \left( |0\rangle_A |0\rangle_B + e^{i\phi} |1\rangle_A |1\rangle_B \right), \\ |\Psi(\phi)\rangle &= \frac{1}{\sqrt{2}} \left( |0\rangle_A |1\rangle_B + e^{i\phi} |1\rangle_A |0\rangle_B \right), \end{aligned}$$

where  $\phi$  is a real number. A popular basis for entangled states is given by the four *Bell states* [Jon12]:

$$\begin{aligned} |\Phi^+\rangle &= |\Phi(0)\rangle = \frac{1}{\sqrt{2}} (|0\rangle_A |0\rangle_B + |1\rangle_A |1\rangle_B), \\ |\Phi^-\rangle &= |\Phi(\pi)\rangle = \frac{1}{\sqrt{2}} (|0\rangle_A |0\rangle_B - |1\rangle_A |1\rangle_B), \\ |\Psi^+\rangle &= |\Psi(0)\rangle = \frac{1}{\sqrt{2}} (|0\rangle_A |1\rangle_B + |1\rangle_A |0\rangle_B), \\ |\Psi^-\rangle &= |\Psi(\pi)\rangle = \frac{1}{\sqrt{2}} (|0\rangle_A |1\rangle_B - |1\rangle_A |0\rangle_B). \end{aligned}$$

### 6.1.2. Quantum measurements

#### Quantum measurement formalism

All the quantum measurements, projective or not, can be considered under the unified framework of generalised measurements. A general measurement is defined in terms of a set of *measurement operators*  $\{\hat{M}_m\}$  which satisfy:

$$\sum_m \hat{M}_m = \mathbb{1}. \quad \text{Completeness relation}$$

When measurement operators are Hermitian ( $\hat{M}_m^\dagger = \hat{M}_m$ ) and orthogonal ( $\hat{M}_m \hat{M}_{m'} = \delta_{m,m'}$ ), they describe a *projective* measurement, and we refer to them as *projectors*. From now on, for the sake of clarity, I will use a symbol  $\hat{P}_m$  for a projector, and  $\hat{M}_m$  of a generalised measurement operator.

Every  $\hat{M}_m$  corresponds to a different measurement outcome  $m$ . A measurement of an initial state  $|\psi\rangle$  gives:

$$\begin{aligned} P(m) &= \langle \psi | \hat{M}_m^\dagger \hat{M}_m | \psi \rangle && \text{Probability of outcome } m \\ |\psi'\rangle &= \frac{\hat{M}_m |\psi\rangle}{\sqrt{P(m)}} && \text{Post-measurement state} \end{aligned} \quad (6.7a)$$

For density operators this generalises to

$$\begin{aligned} P(m) &= \text{Tr}(\hat{M}_m^\dagger \hat{M}_m \rho) && \text{Probability of outcome } m \\ \rho' &= \frac{\hat{M}_m \rho \hat{M}_m^\dagger}{P(m)} && \text{Post-measurement state} \end{aligned} \quad (6.8a)$$

When discussing average measurement outcomes, it is common to describe the measurement in terms of its *observable*. The observable is defined as

$$\hat{M} = \sum_m m \hat{M}_m^\dagger \hat{M}_m \quad (6.9)$$

allowing the average outcomes to be calculated as

$$\langle \hat{M} \rangle = \sum m P(m) = \langle \psi | \hat{M} | \psi \rangle \quad (6.10)$$

for a pure state and

$$\langle \hat{M} \rangle = \text{Tr}(\hat{M} \rho)$$

for any input density matrix. While the expressions above are valid for both projective and non-projective measurement, it is for some reason rare to talk about observables in the context of non-projective measurements. For projective measurements, the observable definition in Eq. (6.9) reduces to  $\hat{M} = \sum_m m P_m$ .

Note that any generalised measurement can be viewed as a projective measurement on a larger system. This is the essence of the Stinespring's dilation theorem [Nie02]. Any generalised measurement  $\hat{M}_m$  may be constructed by preparing an ancilla in state  $|0\rangle$ , coupling it to the system with unitary  $U$ , then measuring the ancilla protectively. The measurement probability  $P(m)$  and the post-measurement state of the system  $\rho'(m)$  obey Equations (6.8a) and (6.7a), with

$$\hat{M}_m = \langle m | U | 0 \rangle.$$

## 6. QUANTUM CONTEXTUALITY

This is known as an *operator-sum decomposition*, and  $\hat{M}_m$  is the *Kraus operator* for the measurement outcome  $m$ . If the measurement outcome is not recorded, the system evolves into

$$\rho' = \sum_m P(m) \rho'(m) = \hat{M}_m \rho \hat{M}_m^\dagger, \quad (6.11)$$

which is a general representation of any quantum process.

### Binary measurements

Trapped ion fluorescence readout is an example of a *binary measurement*. Either fluorescence is detected, and the state is labelled “bright”, or there is no fluorescence, and the state is “dark”. When measuring a qubit, this measurement distinguishes the two basis states  $|0\rangle$  and  $|1\rangle$  from each other, and is ideally projective. A common choice is to associate  $|0\rangle$  with outcome  $m = 1$ , and  $|1\rangle$  with outcome  $m = -1$ . Therefore, the measurement operators are

$$\hat{P}_1 = |0\rangle\langle 0| \quad \text{and} \quad \hat{P}_{-1} = |1\rangle\langle 1|$$

which means the observable (Eq. (6.9)) can be written in one of the following ways

$$\hat{M} = |0\rangle\langle 0| - |1\rangle\langle 1| = 2|0\rangle\langle 0| - \mathbb{1} = \sigma_z$$

where  $\sigma_z$  is the Pauli matrix. This can be extended to a multi-level system (i.e. a qudit) by considering a measurement which distinguishes a bright state  $|0\rangle$  from other states, which are all dark. In this case we can write in general

$$\hat{P}_1 = |0\rangle\langle 0| \quad \text{and} \quad \hat{P}_{-1} = \mathbb{1} - |0\rangle\langle 0| \quad (6.12)$$

which means

$$\hat{M} = 2|0\rangle\langle 0| - \mathbb{1}.$$

These measurement operators show that, if outcome  $+1$  is obtained, the state is projected onto the state  $|0\rangle$ , while if outcome  $-1$  is obtained, the state is projected onto a subspace of states  $\mathbb{1} - |0\rangle\langle 0|$ , preserving coherences present in the initial state [Mal17]. Measurements along arbitrary directions can be performed by combining this measurement with single-qudit unitaries. Consider a general qubit pure state given by

$$|\psi\rangle = U|0\rangle$$

where  $U$  is a single-qudit unitary. Its observable

$$\hat{M}^{|\psi\rangle} = 2|\psi\rangle\langle\psi| - \mathbb{1}. \quad (6.13)$$

satisfies

$$\hat{M}^{|\psi\rangle} = U \hat{M} U^\dagger$$

and therefore<sup>1</sup> can be measured by first applying  $U^\dagger$ , followed by measurement  $\hat{M}$ , and finally applying  $U$ . The final unitary  $U$  is only relevant for the post-measurement state, and does not influence the measurement statistics. Therefore, it only needs to be included in the experiment if there are other measurements following it.

### 6.1.3. Measurement correlators

In both the Bell test and the contextuality tests, we are interested in the correlations between different measurements. For a pair of measurement outcomes  $m, n = \pm 1$  occurring with joint probability  $P(m, n)$ , the average correlator  $m \times n$  is defined as

$$\langle mn \rangle = \sum_{m, n = \pm 1} mn P(m, n).$$

However, the method used to evaluate the correlator is different depending on whether the measurements are bipartite (as in the Bell scenario), or whether they are performed sequentially on a local system (contextual scenario).

#### Bipartite measurements

When the measurement is performed on sub-system  $A$  of a bi-partite system, the total measurement is described by an observable of a form  $\hat{M}_A \otimes \mathbb{1}_B$ , while  $\mathbb{1}_A \otimes \hat{M}_B$  describes a measurement of  $B$ .

In a Bell test, we are interested in the correlation  $(m \times n)$  between the outcome  $m$  of a measurement of  $A$  and the outcome  $n$  of a measurement of  $B$ . These two measurements can be considered as a single quantum measurement with outcome  $(m, n)$  and a projective measurement operator  $\hat{P}_m^{|\psi_1\rangle} \otimes \hat{P}_n^{|\psi_2\rangle}$ , regardless of whether they are performed sequentially or simultaneously. The average correlator between the measurement outcomes performed on a composite pure state  $|\psi_0\rangle$  is given by

$$\langle \hat{M}^{|\psi_1\rangle} \hat{M}^{|\psi_2\rangle} \rangle = \sum_{n, m = \pm 1} mn \langle \psi_0 | \hat{P}_m^{|\psi_1\rangle} \otimes \hat{P}_n^{|\psi_2\rangle} | \psi_0 \rangle = \langle \psi_0 | \hat{M}^{|\psi_1\rangle} \otimes \hat{M}^{|\psi_2\rangle} | \psi_0 \rangle, \quad (6.14)$$

where  $|\psi_1\rangle$  and  $|\psi_2\rangle$  specify the local measurement directions. Because of the tensor product structure,  $\hat{M}^{|\psi_1\rangle} \otimes \hat{M}^{|\psi_2\rangle}$  is a Hermitian operator. Therefore, the measurement correlator can be considered to be an observable itself.

---

<sup>1</sup> Recall that the order of operations in a formula is the reverse of the one in the measurement circuit

## 6. QUANTUM CONTEXTUALITY

### Local sequential measurements

In the contextual scenario, the measurements are performed sequentially on a local system, and can be no longer treated as independent. Suppose that  $\hat{M}^{|\psi_1\rangle}$  is the first observable to be measured, and the initial state is still  $|\psi_0\rangle$ . The average outcome of the first measurement is then (Eq. (6.10)):

$$\langle \hat{M}^{|\psi_1\rangle} \rangle = \langle \psi_0 | \hat{M}^{|\psi_1\rangle} | \psi_0 \rangle \quad (6.15)$$

After the first measurement gives outcome  $m$ , the state is given by (Eq. (6.7a)):

$$|\psi'\rangle = \frac{\hat{P}_m^{|\psi_1\rangle} |\psi_0\rangle}{\sqrt{P(m)}},$$

where  $\hat{P}_m^{|\psi_1\rangle}$  and  $P(m)$  are the projector and outcome probability, respectively. The probability for the second measurement to give outcome  $n$  after the first outcome was  $m$  is

$$P(n | m) = \langle \psi' | \hat{P}_n^{|\psi_2\rangle} | \psi' \rangle = \frac{\langle \psi_0 | \hat{P}_m^{|\psi_1\rangle} \hat{P}_n^{|\psi_2\rangle} \hat{P}_m^{|\psi_1\rangle} | \psi_0 \rangle}{P(m)},$$

and therefore the joint probability is

$$P(n, m) = P(n | m) P(m) = \langle \psi_0 | \hat{P}_m^{|\psi_1\rangle} \hat{P}_n^{|\psi_2\rangle} \hat{P}_m^{|\psi_1\rangle} | \psi_0 \rangle,$$

and the average outcome of the second measurement evaluates to

$$\begin{aligned} \langle \hat{M}^{|\psi_2\rangle} \rangle &= \sum_{n,m=\pm 1} n P(n, m) = \sum_{n,m=\pm 1} n \langle \psi_0 | \hat{P}_m^{|\psi_1\rangle} \hat{P}_n^{|\psi_2\rangle} \hat{P}_m^{|\psi_1\rangle} | \psi_0 \rangle \\ &= \sum_{m=\pm 1} \langle \psi_0 | \hat{P}_m^{|\psi_1\rangle} \left( \sum_{n=\pm 1} n \hat{P}_n^{|\psi_2\rangle} \right) \hat{P}_m^{|\psi_1\rangle} | \psi_0 \rangle \\ &= \sum_{m=\pm 1} \langle \psi_0 | \hat{P}_m^{|\psi_1\rangle} \hat{M}^{|\psi_2\rangle} \hat{P}_m^{|\psi_1\rangle} | \psi_0 \rangle. \end{aligned} \quad (6.16)$$

The sequential measurement correlator can be then evaluated as

$$\langle \hat{M}^{|\psi_1\rangle} \hat{M}^{|\psi_2\rangle} \rangle = \sum_{n,m=\pm 1} mn P(n, m) = \sum_{n,m=\pm 1} mn \langle \psi_0 | \hat{P}_m^{|\psi_1\rangle} \hat{P}_n^{|\psi_2\rangle} \hat{P}_m^{|\psi_1\rangle} | \psi_0 \rangle. \quad (6.17)$$

Plugging in projectors from Eq. (6.12) into Eq. (6.17) and expanding it out, we find that

$$\begin{aligned} \langle \hat{M}^{|\psi_1\rangle} \hat{M}^{|\psi_2\rangle} \rangle &= 1 - 2\|\langle \psi_0 | \psi_1 \rangle\|^2 - 2\|\langle \psi_0 | \psi_2 \rangle\|^2 \\ &\quad + 2\langle \psi_0 | \psi_1 \rangle \langle \psi_1 | \psi_2 \rangle \langle \psi_2 | \psi_0 \rangle \\ &\quad + 2\langle \psi_0 | \psi_2 \rangle \langle \psi_2 | \psi_1 \rangle \langle \psi_1 | \psi_0 \rangle. \end{aligned} \quad (6.18)$$

This formula has a few interesting implications. First of all, the result is unchanged when  $|\psi_1\rangle$  is swapped with  $|\psi_2\rangle$ . Therefore, the measurement correlator is independent of the order in which the measurements are conducted:

$$\langle \hat{M}^{|\psi_1\rangle} \hat{M}^{|\psi_2\rangle} \rangle = \langle \hat{M}^{|\psi_2\rangle} \hat{M}^{|\psi_1\rangle} \rangle.$$

I don't have an intuition about why this happens. Interestingly, Eq. (6.18) can be re-written as

$$\langle \hat{M}^{|\psi_1\rangle} \hat{M}^{|\psi_2\rangle} \rangle = \frac{1}{2} \langle \psi_0 | \{ \hat{M}^{|\psi_1\rangle}, \hat{M}^{|\psi_2\rangle} \} | \psi_0 \rangle \quad (6.19)$$

i.e. half of the anticommutator of two measurements is a valid observable for a correlator of two sequential measurements. Note that Eq. (6.19) reduces to Eq. (6.14) when the measurements commute, and so is actually more general. It turns out this phenomenon exemplifies a general connection between temporal and space-like correlations in quantum physics. This insight can be generalised into a so-called pseudo-density matrix, which treats spatial and temporal correlations within a unified framework [Fit15; Zha17].

### Commuting sequential measurements

In quantum mechanics (QM), measurements do not disturb each other when their observables commute. For binary projective measurements, we find from Eq. (6.13) that

$$[\hat{M}^{|\psi_1\rangle}, \hat{M}^{|\psi_2\rangle}] = 0 \iff [|\psi_1\rangle\langle\psi_1|, |\psi_2\rangle\langle\psi_2|] = 0$$

which happens either in a trivial case of  $|\psi_1\rangle = |\psi_2\rangle$ , or when measurement directions are orthogonal:

$$\langle \psi_1 | \psi_2 \rangle = 0.$$

When these conditions are satisfied, the results for sequential measurements above reduce to

$$\begin{aligned} \langle \hat{M}^{|\psi_2\rangle} \rangle &= 1 - 2\|\langle \psi_0 | \psi_2 \rangle\|^2 = \langle \psi_0 | \hat{M}^{|\psi_2\rangle} | \psi_0 \rangle \\ \langle \hat{M}^{|\psi_1\rangle} \hat{M}^{|\psi_2\rangle} \rangle &= 1 - 2\|\langle \psi_0 | \psi_1 \rangle\|^2 - 2\|\langle \psi_0 | \psi_2 \rangle\|^2 \\ &= \langle \hat{M}^{|\psi_1\rangle} \rangle + \langle \hat{M}^{|\psi_2\rangle} \rangle - 1 \end{aligned} \quad (6.20a)$$

As expected, the statistics of the second measurement are independent of the first measurement direction. Likewise, the measurement correlator is only dependent on measurement statistics of the individual measurements.

## 6.2. Bell theorems

Armed with an understanding of the measurement formalism, and an intuition that non-separability is non-classical, we can now discuss the Bell theorem [Bel64], which attempts to formalise this insight. Consider the so-called CHSH scenario [Cla69] where an agent named Charlie prepares two particles, one of which is given to Alice and the other one to Bob. Alice and Bob perform local measurements on their particles. We use the following notation

- $\lambda$  specifies the full state of the system before the measurements.
- $x$  and  $y$  are measurement settings chosen by Alice and Bob respectively.
- $a$  and  $b$  are measurement outcomes obtained by Alice and Bob respectively.
- $P(ab \mid xy)$  is the probability to measure  $a$  and  $b$  given measurement settings  $x$  and  $y$ .

Note that while  $\lambda$  is usually called a *hidden variable* (HV), it simply represents the state of reality. For example, in QM formalism,  $\lambda$  is the wavefunction of the state prepared by Charlie. In other words,  $\lambda$  represents the state of the system, which may or may not be fully accessible to us. For simplicity, let us assume that Alice and Bob can only choose one of two possible measurement settings and can only obtain one of two possible measurement outcomes.

$$x, y \in \{0, 1\} \quad \text{and} \quad a, b \in \{0, 1\}.$$

We will be interested in the quantity  $p$ , defined as [Cle04]

$$\begin{aligned} 4p = & P(a = b \mid x = 0, y = 0) + P(a = b \mid x = 0, y = 1) \\ & + P(a = b \mid x = 1, y = 0) + P(a \neq b \mid x = 1, y = 1) \end{aligned}$$

which we will write explicitly as:

$$\begin{aligned} 4p = & P(00 \mid 00) + P(11 \mid 00) + P(00 \mid 01) + P(11 \mid 01) \\ & + P(00 \mid 10) + P(11 \mid 10) + P(01 \mid 11) + P(10 \mid 11). \end{aligned} \tag{6.21}$$

We can phrase the problem of maximising  $p$  as a (co-operative) game played by Alice and Bob. If they both choose the measurement setting 1, then score a point if their measurement outcomes disagree, and for all other setting choices, they score a point when their outcomes agree. If all pairs of measurement settings are equally likely, then  $p$  is the probability Alice and Bob will score a point.<sup>2</sup> In order to evaluate  $p$  from Eq. (6.21), we need to be able to evaluate the elements  $P(ab \mid xy)$ . Let's see what classical models have to say about them.

---

<sup>2</sup> In physics, games don't need to be interesting for participants.



### 6.2.1. Bell Theorem #1

The gist of classical attempts to interpret quantum phenomena is to say that quantum randomness could be actually due to unknown and/or unpredictable classical phenomena. In the Bell scenario, that would mean that even if Charlie attempts to send the same state every time, every repetition of the game could in fact correspond to a different  $\lambda$ . Therefore, we write the probabilities in full generality as

$$P(ab | xy) = \sum_{\lambda} P(ab | xy\lambda) P(\lambda | xy). \quad (6.22)$$

Now we make the *free-will assumption* [Ara16]

$$P(\lambda | xy) = P(\lambda) \quad (6.23)$$

Eq. (6.23) does not actually say that Alice and Bob have free will<sup>3</sup>, but rather that measurement settings  $x, y$  can be chosen independently of the state of the measured system. As far as I'm aware, it's implicit in all theories, classical or quantum. Substituting Eq. (6.23) into Eq. (6.22) gives

$$P(ab | xy) = \sum_{\lambda} P(ab | xy\lambda) P(\lambda). \quad (6.24)$$

The next assumption is *realism*<sup>4</sup> or *determinism*

$$P(ab | xy\lambda) = 0 \quad \text{or} \quad P(ab | xy\lambda) = 1. \quad (6.25)$$

In other words, the state of reality  $\lambda$  completely determines the measurement outcome. This would match most classical intuitions, and be at odds with the formalism of QM, which says that measurement outcomes are fundamentally unpredictable. Eq. (6.25) implies algebraically that

$$P(ab | xy\lambda) = P(a | xy\lambda) P(b | xy\lambda). \quad (6.26)$$

The final assumption is *locality*, which implies

$$P(a | xy\lambda) = P(a | x\lambda) \quad \text{and} \quad P(b | xy\lambda) = P(b | y\lambda). \quad (6.27)$$

This assumption, motivated by the theory of relativity, states that, for any  $\lambda$  the choice of Bob's measurement settings does not influence Alice's measurement outcomes, and vice

---

<sup>3</sup> That assumption would not be compatible with classical physics.

<sup>4</sup> The naming of this assumption is somewhat contentious - I think nowadays people are so used to QM that they could believe in a non-deterministic reality.

## 6. QUANTUM CONTEXTUALITY

versa. Equations (6.27) and (6.26) together imply factorisability [Fin82]:

$$P(ab | xy\lambda) = P(a | x\lambda) P(b | y\lambda). \quad (6.28)$$

Applying factorisability to Eq. (6.24) we obtain

$$P(ab | xy) = \sum_{\lambda} P(a | x\lambda) P(b | y\lambda) P(\lambda). \quad (6.29)$$

In order to constrain  $p$ , we only need to consider 16 possible values of  $\lambda$ . To see this, observe that there are 4 possible truth tables for  $P(a | x)$ , since in order to define a truth table, we need to make two independent choices

1. Either  $P(0 | 0) = 1$  and  $P(1 | 0) = 0$ , or  $P(0 | 0) = 0$  and  $P(1 | 0) = 1$ .
2. Either  $P(0 | 1) = 1$  and  $P(1 | 1) = 0$ , or  $P(0 | 1) = 0$  and  $P(1 | 1) = 1$ .

Since there are 4 possible truth tables for  $P(a | x)$ , and 4 possible truth tables for  $P(b | y)$ , then there are 16 possible truth tables for  $P(ab | xy)$ . In order to construct a classical bound for  $p$ , we simply need to check those 16 cases. The bound is obtained by only considering  $\lambda$  which maximises  $p$ . By iterating through all the cases, we find that one of the scenarios which maximises  $p$  is

$$\begin{aligned} P(a = 0 | x = 0) = 1 & \quad \text{and} \quad P(a = 1 | x = 1) = 1 \\ P(b = 0 | x = 0) = 1 & \quad \text{and} \quad P(b = 0 | x = 1) = 1 \end{aligned}$$

Resulting in a classical bound on  $p$

$$p \leq p_L = \frac{3}{4}. \quad (6.30)$$

This is known as the CHSH inequality [Cla69]. We know we this bound is exceeded in the real world. Specifically, consider the situation where Charlie prepares a maximally entangled state  $|\Phi^+\rangle$ . Following [Cle04], define:

$$\begin{aligned} |\phi_0(\theta)\rangle &= \cos(\theta)|0\rangle + \sin(\theta)|1\rangle \\ |\phi_1(\theta)\rangle &= -\sin(\theta)|0\rangle + \cos(\theta)|1\rangle. \end{aligned}$$

which describe orthogonal measurement directions for a single qubit. For measurement settings  $x$  and  $y$ , Alice and Bob will measure operators  $X_x^a$  and  $Y_y^b$  given by:

$$\begin{aligned} X_0^a &= |\phi_a(0)\rangle\langle\phi_a(0)| \\ X_1^a &= |\phi_a(\pi/4)\rangle\langle\phi_a(\pi/4)| \\ Y_0^b &= |\phi_b(\pi/8)\rangle\langle\phi_b(\pi/8)| \\ Y_1^b &= |\phi_b(-\pi/8)\rangle\langle\phi_b(-\pi/8)| \end{aligned} \quad (6.32a)$$

for  $(a, b) \in \{0, 1\}$ . From Eq. (6.14) we find:

$$P(ab | xy) = \langle \Phi^+ | X_x^a Y_y^b | \Phi^+ \rangle$$

With this formula, we find that every element on the right hand side in Eq. (6.21) equals  $\frac{1}{2} \cos^2(\pi/8)$ , and therefore in total

$$p_{\text{QM}} = \cos^2(\pi/8) = \frac{2 + \sqrt{2}}{4} \approx 0.85. \quad (6.33)$$

This is the *Tsirelson's Bound* for a Bell inequality violation [Cir80], and the maximum value that  $p$  can achieve in QM. The fact that  $p_{\text{QM}} > p_L$  proves that QM is inconsistent with any model which assumes free will, determinism (realism) and locality.

### 6.2.2. Extended Bell's Theorem and the role of determinism

A conclusion that most physicists draw from Bell's theorem is that the assumption of determinism (or realism) needs to be abandoned - this seems like the easiest way out. However, it turns out that it is not enough [Cla74]. In this section, we discuss a follow-up to Bell's theorem, which shows that that local *stochastic* models are also bounded by  $p_{\text{max}}$ .

To see how that comes about, let us analyse the problem without assuming determinism. Let us also write the assumption of locality in a different way

$$P(a | xy\lambda b) = P(a | x\lambda) \quad \text{and} \quad P(b | xy\lambda a) = P(b | y\lambda). \quad (6.34)$$

Eq. (6.34) looks very similar to Eq. (6.27), but is actually a stronger assumption, known as *local causality* [Bel85]. The idea of local causality is that  $P(a)$  is independent of everything outside of Alice's light cone, which includes both Bob's measurement setting  $y$  and Bob's measurement outcome  $b$ . With Eq. (6.34), we can derive Bell inequality without assuming Eq. (6.25) [Ara]. This starts by using Bayes rule:

$$P(ab | xy\lambda) = P(a | xy\lambda b) P(b | xy\lambda)$$

which, using Eq. (6.34), reduces to

$$P(ab | xy\lambda) = P(a | x\lambda) P(b | y\lambda) \quad (6.35)$$

which is the factorisability condition from Eq. (6.28) [Fin82]! This immediately implies that  $P(ab | xy)$  is still given by Eq. (6.29) and therefore the classical bound is still given by  $p_L$ <sup>5</sup>.

---

<sup>5</sup> To be precise, to prove this, we need to consider all possible stochastic truth tables for  $P(a | x)$  and  $P(b | y)$ , while before we only had to consider deterministic truth tables, filled with 0's and 1's. However, all probabilistic truth tables can be written as a weighted sum of deterministic truth tables. So in fact the maximum value of  $p$  when summing over all probabilistic truth tables is the same as when summing over all deterministic truth tables

Therefore, the local deterministic bound on  $p$  in Eq. (6.30) is also a locally causal bound. The extended Bell's theorem implies that abandoning determinism does not allow one to completely rescue locality. If one believes in QM and free will, it is possible to believe in locality in the sense of Eq. (6.27), but not in locality in the sense of local causality in Eq. (6.34).

I find the extended Bell's theorem fascinating and important. On a philosophical level, it makes it clear that quantum correlations are not purely due to non-determinism, and that some kind of non-locality is at play<sup>6</sup>. While the original theorem rejects one kind of models – deterministic ones – the extension allows rejecting all kinds of probabilistic models that respect local causality.

Aside from their fundamental importance, Bell inequalities are the backbone of many proposed technologies. Most excitingly, recent years have brought about the development of device independent (DI) quantum random number generation (QRNG) [Pir14] and DI quantum key distribution (QKD) protocols [Ací07]. Those protocols are secure against attacks by adversaries limited only by QM.

### 6.2.3. Non-signaling models

The conclusion of Bell's theorem is that correlations present in QM exceed those available to local deterministic models. The conclusion of the extended Bell's theorem is that they also exceed those available to locally causal models. This is the sense in which QM can be considered non-local. In this section, we attempt to quantify exactly how non-local it is.

According to the current laws of physics, information cannot be transferred faster than the speed of light. For Alice and Bob, this implies that they cannot influence each other's measurement statistics simply by changing their measurement directions. This assumption is called *non-signaling* (NS) and can be written as [Mas06]:

$$P(a | xy) = P(a | x) \quad \text{and} \quad P(b | xy) = P(b | y). \quad (6.36)$$

Eq. (6.36) looks very similar to the locality assumption (Eq. (6.27)), but it's very much not the same. In fact, locality implies NS, but not vice versa. While local deterministic models (and locally causal models) are factorisable, NS models may not be, as we shall see below.

QM itself is a non-deterministic NS model. Therefore, an interesting question is whether the NS conditions in Eq. (6.36) are what limits the maximum value of  $p$  to  $p_{\text{QM}}$ . The answer turns out to be negative, and it has some important implications. To see this, let us consider a general NS model. Following [Col09], we write a general truth table for  $P(ab | xy)$  in the way shown in Tab. 6.1.

The advantage of writing it this way is that the NS condition from Eq. (6.36) is satisfied if, for every row and column, the sum of the first two entries equals the sum of the second

---

<sup>6</sup> In fact, it even hints that non-determinism is not even a particularly important factor.

**Table 6.1:** A general truth table for  $P(ab | xy)$ 

		y		0		1	
		b		0	1	0	1
x	a						
0	0	P(00   00)		P(01   00)		P(00   01)	
	1	P(10   00)		P(11   00)		P(10   01)	
1	0	P(00   10)		P(01   10)		P(00   11)	
	1	P(10   10)		P(11   10)		P(10   11)	

two entries. For example, for the first row, this condition reads

$$P(00 | 00) + P(01 | 00) = P(00 | 01) + P(01 | 01).$$

which is equivalent to

$$P(a = 0 | 00) = P(a = 0 | 01)$$

and hence

$$P(a = 0 | x = 0, y) = P(a = 0 | x = 0)$$

which is exactly the NS condition for  $a = 0$  and  $x = 0$ . Other rows give other NS conditions for  $P(a | x)$ , and the columns give NS conditions for  $P(b | y)$ . Such truth table corresponds to either a sum over deterministic theories with different  $\lambda$ , or a non-deterministic model with a single  $\lambda$ .

We are now in a position to see a truth table for a NS model which exceeds the quantum bound. This is known as a Popescu-Rohrlich (PR) box [Pop94] and is given in Table 6.2.

**Table 6.2:** Truth table for a NS model with maximal  $p$ 

		y		0		1	
		b		0	1	0	1
x	a						
0	0	1/2		0	1/2	0	0
	1	0		1/2	0	1/2	0
1	0	1/2		0	0	1/2	1/2
	1	0		1/2	1/2	0	0

It is easy to check that the model in Table 6.2 is NS. However, it is not factorisable, since it cannot be decomposed onto separate truth tables for  $P(a | x)$  and  $P(b | y)$ , as Eq. (6.35) would have it. Therefore, this is not a truth table of a locally causal model. Inserting the

## 6. QUANTUM CONTEXTUALITY

values from Table 6.2 into Eq. (6.21) we obtain

$$p = p_{\text{NS}} = 1 > p_{\text{QM}} > p_L.$$

So in fact, general NS theories can violate Bell inequalities more than QM. So in this case, QM seems to be more non-local than classical intuitions allow, but less non-local than special relativity actually permits!

### 6.2.4. Nonlocal content and Chained Bell Inequalities

In the case of Bell inequalities, QM gives correlations that are stronger than any deterministic NS theory, but weaker than some probabilistic NS theories. An interesting way of looking at the relation between NS and QM is using *chained Bell inequalities*. We start by considering a different figure of merit for a Bell test, defined as

$$I_2 = P(a = b \mid 10) + P(a \neq b \mid 00) + P(a \neq b \mid 01) + P(a \neq b \mid 11)$$

For this witness, following identical steps as above, we can derive the following bounds [Col09]:

$$\begin{array}{ll} I_2 \geq I_{2,L} = 1 & \text{in local models} \\ I_2 \geq I_{2,T} = 4 \sin^2(\pi/8) = 2 - \sqrt{2} \approx 0.59 & \text{in QM} \\ I_2 \geq I_{2,\text{NS}} = 0 & \text{in NS models} \end{array}$$

These are known as *chained Bell inequalities* [Bra90]. Note that here the word ‘local’ is a shortcut for “local deterministic or causally local”. As before, there is a significant gap between the quantum limit  $I_{2,T}$  and the NS limit  $I_{2,\text{NS}}$ . We can now extend the scenario considered so far by allowing Alice and Bob to make one of  $N > 2$  different measurements each,  $(x, y) \in \{0, \dots, N-1\}$ , with binary outcomes  $(a, b) \in \{0, 1\}$ . Consider the following witness:

$$\begin{aligned} I_N = & P(a = b \mid N-1, 0) + P(a \neq b \mid 00) + P(a \neq b \mid 01) + P(a \neq b \mid 11) \\ & + P(a \neq b \mid 12) + \dots + P(a \neq b \mid N-1, N-1), \end{aligned}$$

which reduces to  $I_2$  for  $N = 2$ . In order to bound  $I_N$ , we first note that the classical and quantum bounds are unchanged. The optimal quantum strategy is obtained by changing the measurement angles in Eq. (6.32a) from  $\pi/4$  to  $\pi/(2N)$ . All in all, the bounds become

$$\begin{array}{ll} I_N \geq I_L = 1 & \text{in local models} \\ I_N \geq I_{N,T} = 2N \sin^2\left(\frac{\pi}{4N}\right) & \text{in QM} \\ I_N \geq I_{N,\text{NS}} = 0 & \text{in NS models} \end{array}$$

The interesting observation is that the gap between NS models and QM vanishes in the limit of large  $N$ . One way to formalise it is by defining the *nonlocal content*  $q$  [Chr15]. The idea is to model the observed correlations as a sum of local and non-local behaviours [Sca19]. Assuming NS, any observed value of  $I_N$  can be explained by assuming Alice and Bob behave non-locally with probability  $q$ , and locally with probability  $1 - q$ , where

$$I_N = (1 - q)I_L + qI_{NS}$$

which implies that

$$q = 1 - I_N \tag{6.39}$$

Indeed, if QM is correct, then as  $N \rightarrow \infty$ , the system becomes fully non-local ( $q \rightarrow 1$ ). This might sound like an odd way to phrase it, but this framing of the problem has allowed for many interesting findings and applications:

- [Col11] has used the fact that  $q \rightarrow 1$  to show that no extension of quantum theory could have improved predictive power. Crucially, added predictive power of any post-quantum theory can be bounded by experimental measurements of  $q$ .
- There are QKD protocols based on Chained Bell Inequalities which are resistant to eavesdropping by adversaries bound only by NS, not just by QM [Bar05; Ací06b]. In absence of noise,  $q \rightarrow 1$  allows for most efficient information transfer, even for non-local adversaries [Ací06a].
- There are DI QRNG protocols that allow certifiable randomness generation. An adversary bounded only by NS can guess the random number with a probability of at most  $1 - q/2$  [Sca19].

While I'm not terribly worried about NSA spying on people with non-local NS methods<sup>7</sup>, I think that it's very interesting to discover how non-local real-world correlations can be. As of today, the largest measured non-local content for a chained Bell scenario was  $q = 0.874(1)$ , measured with  $N = 18$  settings [Chr15].

### 6.2.5. Bell inequalities in correlator formalism

So far we assumed that measurement outcomes take values  $(a, b) \in \{0, 1\}$ . A different common way of writing down Bell inequalities is in terms out correlators of measurements  $(X, Y)$  with outcomes  $(A, B) \in \{1, -1\}$

$$\langle XY \rangle = P(A = B \mid XY) - P(A \neq B \mid XY)$$

---

<sup>7</sup> If they are this far ahead of mainstream science, they can probably break NS too

## 6. QUANTUM CONTEXTUALITY

For the sake of completeness, below are the main results from sections above, re-expressed in correlator formalism. The CHSH scenario considers the witness

$$E = \langle X_1 Y_1 \rangle + \langle X_1 Y_2 \rangle + \langle X_2 Y_1 \rangle - \langle X_2 Y_2 \rangle$$

and states that

$$\begin{aligned} E &\leq E_L = 2 && \text{in local models} \\ E &\leq E_{\text{QM}} = 2 + \sqrt{2} && \text{in QM} \\ E &\leq E_{\text{NS}} = 4 && \text{in NS models} \end{aligned}$$

The chained Bell inequalities extend this witness to

$$E_N = \sum_{i=1}^N \langle X_i Y_i \rangle + \sum_{i=1}^{N-1} \langle X_{i+1} Y_i \rangle - \langle X_1 Y_N \rangle \quad (6.41)$$

and state that

$$\begin{aligned} E_N &\leq E_{N,L} = 2N - 2 && \text{in local models} \\ E_N &\leq E_{N,\text{QM}} = 2N \cos\left(\frac{\pi}{2N}\right) && \text{in QM} \\ E_N &\leq E_{N,\text{NS}} = 2N && \text{in NS models} \end{aligned}$$

The nonlocal content  $q$  can be defined via

$$(1 - q)E_{N,L} + qE_{N,\text{NS}} = E_{N,\text{QM}}$$

and leads to the same expression as Eq. (6.39)

$$q = 1 - 2N \sin^2\left(\frac{\pi}{4N}\right)$$

### 6.3. Quantum contextuality and KCBS test

Because Bell inequalities apply only to non-local systems, it is hard to believe they capture the whole non-classicality of QM. For example, quantum advantage in sensing and computation arises already in a local setting. Therefore, some non-classical correlations should be observable already for local systems. Quantum contextuality attempts to generalise the concepts introduced above without referring to the notion of locality. The onset of contextual correlations is known in certain scenarios to coincide with the possibility of universal quantum computation [How14].

Consider a general scenario where measurements  $(x, y)$  are performed on some system, with measurement outcomes  $(a, b)$ . In QM, measurements do not always commute. That



implies that, in general

$$P(b \mid xy) \neq P(b \mid y),$$

so the outcome of measurement  $y$  might depend on what other measurement  $x$  was performed. However, we will consider the special case of *compatible measurements*. Measurements  $x$  and  $y$  are defined as compatible when

$$P(a \mid xy) = P(a \mid x) \quad \text{and} \quad P(b \mid xy) = P(b \mid y). \quad (6.43)$$

The meaning of conditional probability in this Eq. (6.43) is that the  $P(a \mid x)$  is independent of whether  $x$  is measured alone or together with  $y$ . It also generalises to all possible combinations of measurement settings. For example, if measurements  $x, y, x$  are performed in a sequence, then  $P(aba \mid xyx) = P(aa \mid xx)$ , etc [Cab09]. We can immediately see that the compatibility equations (6.43) are the same as NS Equations (6.36). While NS assumption only makes sense for space-like separated systems, compatibility can be applied more broadly. For example, in QM, any projective measurements whose operators commute are compatible.

Now consider a general model that generates  $P(ab \mid xy)$ , which is still given by Eq. (6.22). Like before, we make the free-will assumption (Eq. (6.23)) and the determinism assumption (Eq. (6.25)), leading us once again to

$$P(ab \mid xy) = \sum_{\lambda} P(a \mid xy\lambda) P(b \mid xy\lambda) P(\lambda).$$


Now we cannot invoke locality anymore, but this time we invoke the assumption of *non-contextuality* (NC) [Cab10b]. It states that if two measurements  $(x, y)$  are compatible, then they are compatible for every state of reality  $\lambda$ :

$$P(a \mid xy) = P(a \mid x) \implies P(a \mid xy\lambda) = P(a \mid x\lambda). \quad (6.44)$$

NC essentially says that if two measurements  $x, y$  never influence each other's statistics (are compatible), then the underlying theory does not include terms where  $x$  and  $y$  influence each other's statistics (are incompatible). I devote Sec. 6.6.2 and Sec. 6.6.3 to assessing the sensibility of this assumption (spoiler: it is sensible, but not nearly as bulletproof as locality).

It is clear that NC allows us to make the same step as locality did, that is, if  $(x, y)$  are compatible, then:

$$P(a \mid xy\lambda) = P(a \mid x\lambda) \quad \text{and} \quad P(b \mid xy\lambda) = P(b \mid y\lambda).$$



**Figure 6.1:** Compatibility structure of KCBS scenario. Measurement setting  $x_i$  is compatible with measurement settings  $x_{i\pm 1}$ , with sum taken modulo 5.

and we once again can write that

$$P(ab | xy) = \sum_{\lambda} P(a | x\lambda) P(b | y\lambda) P(\lambda). \quad (6.45)$$

This means that NC inequalities can be constructed by maximising over all deterministic truth tables for  $P(a | x)$  and  $P(b | y)$ .

### 6.3.1. KCBS inequality

The idea behind NC inequalities is to exploit the fact that, in QM, compatibility is not transitive [Mer93]. That is, if measurement  $y$  is compatible with measurements  $x$  and  $z$  individually, it does not mean that  $x$  is compatible with  $z$ . This non-classical phenomenon allows differentiating quantum and classical predictions. The simplest NC inequality for a local system is the Klyachko-Can-Biniciöglu-Shumovsky (KCBS) inequality with 5 observables [Kly08]. In QM, this inequality can be violated with a single 3-level system (qutrit), which is also the smallest dimension where contextuality can be observed [Cab10b].

Consider 5 measurements  $(x_1, x_2, x_3, x_4, x_5)$  with measurement outcomes  $(A_1, A_2, A_3, A_4, A_5)$ , with  $A_i \in \{-1, 1\}$ . We assume measurements are pairwise compatible:  $x_i$  is compatible with  $x_{i+1}$ , with the sum taken modulo 5. This structure can be shown as a graph in Fig. 6.1, where measurements connected with a solid line are compatible.

Consider the following quantity

$$S = \sum_{i=1}^5 \langle A_i A_{i+1} \rangle. \quad (6.46)$$

This sum can be expressed in terms of probabilities of measurement outcomes using

$$\begin{aligned} \langle A_i A_{i+1} \rangle &= P(A_i = 1, A_{i+1} = 1) + P(A_i = -1, A_{i+1} = -1) \\ &\quad - P(A_i = -1, A_{i+1} = 1) - P(A_i = 1, A_{i+1} = -1). \end{aligned}$$

Since measurements in this expression are compatible, it makes sense to consider predictions of a deterministic NC model. Using Equation (6.45), we can factorise every probability in the sum as

$$P(A_i = \pm 1, A_{i+1} = \pm 1) = P(A_i = \pm 1, \lambda) P(A_{i+1} = \pm 1, \lambda) P(\lambda),$$

which allows us to write

$$\langle A_i A_{i+1} \rangle = \sum_{\lambda} (P(A_i = 1, \lambda) - P(A_i = -1, \lambda)) (P(A_{i+1} = 1, \lambda) - P(A_{i+1} = -1, \lambda)) P(\lambda).$$

Of course  $P(A_i = 1, \lambda) - P(A_i = -1, \lambda)$  is nothing but the expectation value of the measurement of  $A_i$  for a given  $\lambda$ . Consequently, we can simply write

$$\langle A_i A_{i+1} \rangle = \sum_{\lambda} \langle A_i, \lambda \rangle \langle A_{i+1}, \lambda \rangle P(\lambda).$$

Note that since the model is deterministic, either  $P(A_i = 1, \lambda) = 1$  and  $P(A_i = -1, \lambda) = 0$  or  $P(A_i = 1, \lambda) = 0$  and  $P(A_i = -1, \lambda) = 1$ . Therefore, while  $\langle A_i A_{i+1} \rangle$  is a real number in the range  $(-1, +1)$ ,  $\langle A_i, \lambda \rangle$  is either  $+1$  or  $-1$ . Hence to find the possible values of  $S$ , it suffices to test try all  $2^5$  possible values of  $(\langle A_i, \lambda \rangle, \langle A_2, \lambda \rangle, \dots, \langle A_5, \lambda \rangle) = (\pm 1, \pm 1, \dots, \pm 1)$ .

By considering all possible assignments, we find that  $S \geq -3$ . This bound can be understood as follows. The value  $S = -3$  is achievable, for example for  $(\langle A_i, \lambda \rangle, \langle A_2, \lambda \rangle, \langle A_3, \lambda \rangle, \langle A_4, \lambda \rangle, \langle A_5, \lambda \rangle) = (+1, -1, +1, -1, +1)$ . The next lower algebraically allowed value would be  $S = -5$ . This could happen if and only if, for every  $i$

$$\langle A_i, \lambda \rangle = -\langle A_{i+1}, \lambda \rangle,$$

with the sum taken modulo 5. But that would imply that

$$\langle A_1, \lambda \rangle = -\langle A_{i+2}, \lambda \rangle = \langle A_3, \lambda \rangle = -\langle A_4, \lambda \rangle = \langle A_5, \lambda \rangle = -\langle A_1, \lambda \rangle,$$

which is a contradiction. Therefore we derived the KCBS inequality:

$$S \geq S_{\text{NC}} = -3, \quad (6.47)$$

which again holds for all deterministic NC models. This will see in the next section that it is violated in QM, with

$$S \geq S_{\text{QM}} = 5 - 4p_{\text{QM}} = 5 - 4\sqrt{5} \approx -3.944, \quad (6.48)$$

Interestingly, the original KCBS paper makes a stronger claim, suggesting that Eq. (6.47) “provides a test for arbitrary hidden variables model, context free or not” [Kly08]. Bizarrely, none of the many papers that quote this result comment on this statement, but none seem to agree with it either. So let me make it clear that the quoted statement is incorrect, and that assumption of NC is fundamental to their derivation.

### 6.3.2. Qutrit measurements and quantum violation

Let us now construct a sequence of QM measurements which violates the KCBS inequality. I could just state them, like I did in Sec. 6.2 for the Bell inequality, but it is instructive to see how one may come up with them. Consider a 3-level Hilbert space with basis states  $|0\rangle, |1\rangle$  and  $|2\rangle$ . A general pure state of a qutrit is given by:

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle + \gamma|2\rangle$$

where  $|\alpha|^2 + |\beta|^2 + |\gamma|^2 = 1$ . It will be sufficient to only consider states with  $(\alpha, \beta, \gamma) \in \mathbb{R}$ . This way, every  $|\psi\rangle$  can be visualised as a vector on a unit sphere in a 3D space, with co-ordinates  $(\alpha, \beta, \gamma)$ . We will call this space a *qutrit sphere*.


Let us now define binary measurements in this space. Following Eq. 6.13, the observable for a measurement along  $|\psi\rangle$  is defined as:

$$M = 2|\psi\rangle\langle\psi| - \mathbb{1}$$

In order to re-create the compatibility structure of KCBS inequality, we need to find five measurements  $M_i$ , each of which is compatible with  $M_{i\pm 1}$ . Recall from Section 6.1.3 that this is equivalent to finding measurement directions  $|\psi_i\rangle$  which satisfy pairwise orthogonality conditions:

$$\langle\psi_i|\psi_{i\pm 1}\rangle = 0$$

Consider the construction as shown in Figure 6.2. We will associate  $|0\rangle$  with  $z$ ,  $|1\rangle$  with  $y$  and  $|2\rangle$  with  $x$ . The first measurement direction  $|\psi_1\rangle$  is in the  $zy$  plane and makes an angle  $\theta$  to  $z$ . The other measurement directions are lying on a regular pentagon in the  $xy$  plane.



**Figure 6.2:** Measurement directions for KCBS inequality on a qutrit sphere. States connected by dashed lines correspond to compatible measurements when  $\theta = \theta_5$ , when their vectors are orthogonal. State  $|\psi_1\rangle$  is obtained rotating  $|0\rangle$  around  $y$  axis by  $\theta$ . State  $|\psi_{i+1}\rangle$  is obtained rotating  $|\psi_i\rangle$  around  $z$  axis by  $\frac{4\pi}{5}$ . Figure from [Mal17] ©2017 American Physical Society

Before doing any maths, consider what happens to the angle between  $|\psi_0\rangle$  and  $|\psi_1\rangle$  as  $\theta$  is varied from 0 to  $\pi/2$ . For  $\theta = 0$ , the angle between them is 0. For  $\theta = \pi/2$ , all the state vectors are in the same plane, and the angle is  $4\pi/5$ . Therefore, from continuity, there is some angle  $0 < \theta_5 < \pi/2$  for which  $\langle \psi_2 | \psi_1 \rangle = 0$  and the corresponding measurements are compatible. By symmetry, at this point,  $\langle \psi_{i+1} | \psi_i \rangle = 0$  for all  $i$ . Consequently, for  $\theta = \theta_5$ , all the measurements along direction connected by dashed lines in Figure are compatible, and so the measurements reproduce the compatibility structure of the KCBS inequality. To be more quantitative, we can define the five states using rotations  $R_y$  and  $R_z$  around  $y$  and  $z$  axes respectively, as shown in Figure 6.2. The first measurement direction is given by  $|\psi_1\rangle = U_1|0\rangle$ , where

$$U_1 = R_y(\theta).$$

The angle between subsequent vertices of a regular pentagon is  $2\pi/5$ . The second direction is given by rotating  $|\psi_1\rangle$  around  $z$  this angle twice:

$$|\psi_2\rangle = R_z^2(2\pi/5)|\psi_1\rangle = R_z^2(2\pi/5)R_y(\theta)|0\rangle.$$

By this construction, a general measurement direction  $|\psi_i\rangle$  is given by

$$|\psi_i\rangle = U_i|0\rangle \quad \text{where} \quad U_i = R_z^{2i-2}(2\pi/5)R_y(\theta) \quad (6.49)$$

## 6. QUANTUM CONTEXTUALITY

Note that the rotations in Eq. (6.49) are given in Cartesian space. The rotation matrices are explicitly stated in Sec. 7.4.1. Using this definition, we find that

$$\langle \psi_i | \psi_{i+1} \rangle = \cos^2 \theta - \frac{1}{4}(1 + \sqrt{5}) \sin^2 \theta.$$

This gives the orthogonality condition

$$\theta = \theta_5 = \arctan \left( \frac{2}{\sqrt{1 + \sqrt{5}}} \right) \approx 48^\circ. \quad (6.50)$$

In order to evaluate the KCBS witness in Eq. (6.46), we need to calculate expectation values of correlators. For an initial state  $|0\rangle$ , and assuming compatible measurements, Eq. (6.20a) gives

$$\langle A_i A_{i+1} \rangle = 1 - 2\|\langle \psi_0 | \psi_1 \rangle\|^2 - 2\|\langle \psi_0 | \psi_2 \rangle\|^2 = 1 - 4\cos^2(\theta_5). \quad (6.51)$$

where we used the fact the every state  $|\psi_i\rangle$  has the same overlap with the initial state, given by

$$\langle 0 | \psi_i \rangle = \cos(\theta_5).$$

Plugging in  $\theta_5$  from Eq. (6.50) gives:

$$\langle A_i A_{i+1} \rangle = 1 - \frac{4}{\sqrt{5}} \approx -0.789$$

Thus the whole KCBS witness in Eq. (6.46) evaluates to

$$S = \sum_{i=1}^5 \langle A_i A_{i+1} \rangle = S_{\text{QM}} = 5 - 4\sqrt{5} \approx -3.944$$

which significantly violates the NC bound of  $S_{\text{NC}} = -3$  in the KCBS inequality. This is also the maximal violation of this inequality in QM, meaning that  $S \geq S_{\text{QM}}$  for other initial states or for other sets of five *compatible* measurements. It does not, however, mean that  $S \geq S_{\text{QM}}$  for any set of five projective measurements in QM. For example, as we will see later, if some of the measurement directions  $|\psi_i\rangle$  are mis-calibrated, this can lead to  $S < S_{\text{QM}}$ , limited only by the algebraic limit  $S \geq -5$ . However, mis-calibrated measurements will

also, in general, be incompatible, violating the assumption in Eq. (6.45). To summarise:

$$\begin{array}{ll} S \geq S_{\text{NC}} = -3 & \text{for deterministic NC models} \\ S \geq S_{\text{QM}} \approx -3.944 & \text{for compatible measurements in QM} \\ S \geq -5 & \text{for incompatible measurements} \end{array}$$

We also note that a general model for compatible measurements, not bounded by determinism or NC, is also only bounded by the algebraic limit:

$$S \geq S_{\text{NS}} = -5 \quad \text{for compatible measurements}$$

We call this bound  $S_{\text{NS}}$  since compatibility conditions in Eq. (6.43) are mathematically equivalent to NS conditions in Eq. (6.36). In summary, the methods of this section allow us to construct an inequality which is satisfied for deterministic NC models, but can be violated by measuring correlators of five different pairwise-compatible measurements in on a single qutrit.

### 6.3.3. Exclusivity and connection to Bell inequality

KCBS inequality can be re-cast in terms of measurements with outcomes  $(0, 1)$ . The witness then reads

$$p = \sum_{i=1}^5 P(0, 1 \mid i, i+1) \quad (6.53)$$

which we write in full generality as

$$p = \sum_{\lambda} \sum_{i=1}^5 P(0, 1 \mid i, i+1, \lambda) P(\lambda) \quad (6.54)$$

In Eq. (6.53), all measurements are compatible, since every measurement  $x_i$  is performed only either with  $x_{i-1}$  or with  $x_{i+1}$ . Assuming determinism and NC we can therefore write

$$P(0, 1 \mid i, i+1, \lambda) = P(0 \mid i, \lambda) P(1 \mid i+1, \lambda)$$

## 6. QUANTUM CONTEXTUALITY

Again, we need to consider all possible truth tables for these measurements. One way to simplify the brute-force approach for bounding  $p$  is as follows. Note that, for any  $\lambda$ :

$$\begin{aligned}
& P(0, 1 \mid i, i+1, \lambda) = 1 \\
& \implies P(0 \mid i, \lambda) P(1 \mid i+1, \lambda) = 1 \\
& \implies P(0 \mid i, \lambda) = 1 \quad \text{and} \quad P(1 \mid i+1, \lambda) = 1 \\
& \implies P(1 \mid i, \lambda) = 0 \quad \text{and} \quad P(0 \mid i+1, \lambda) = 0 \\
& \implies P(0, 1 \mid i-1, i, \lambda) = 0 \quad \text{and} \quad P(0, 1 \mid i+1, i+2, \lambda) = 0
\end{aligned}$$

Therefore, if the element  $i$  of the inner sum in Eq. (6.54) is equal to 1, then elements  $i-1$  and  $i+1$  must be equal to 0. There are five elements in the inner sum, so at most two of them can be equal to 1. Therefore

$$\sum_{i=1}^5 P(0, 1 \mid i, i+1, \lambda) \leq 2$$

which is an alternative form of the KCBS inequality

$$p \leq p_{\text{NC}} = 2$$

that is satisfied by all deterministic NC models. The QM bound is


$$p \leq p_{\text{QM}} = \sqrt{5}$$

The key step of the proof was that, given the compatibility structure in Figure 6.1, determinism and NC imply that, for any two probabilities connected by a line, if one of them is equal to 1, the other one is equal to 0. In other words, determinism means that those probabilities correspond to discrete events, and NC implies that these events are *exclusive*, i.e. if one of them takes place, the other one does not [Cab13]. Hence, Fig. 6.3 displays the *exclusivity graph* of the KCBS scenario.

We can now make a connection between  $p_{\text{NC}}$  and graph theory.  $p_{\text{NC}}$  is equal to the maximum number of vertices of the graph  $G$  in Fig. 6.3 which can be simultaneously equal to 1. And since vertices connected by an edge cannot be both equal to 1,  $p_{\text{NC}}$  is equal to the size of the largest sub-graph of  $G$  where no two elements are directly connected. This sub-graph is known as the *maximal independent set*, its size is called the *independence number* and is denoted as  $\alpha(G)$ . For KCBS inequality,  $\alpha(G) = 2$  and therefore  $p_{\text{NC}} = 2$ . One of the maximal subgraphs of  $G$  is shown in Figure 6.3 as red dots.

The exclusivity framework allows us to make more powerful observations. It turns out that, given an exclusivity graph  $G$ , the maximum quantum value of  $p$  is bounded by a graph-theoretic quantity called the *Lovasz number*  $\vartheta(G)$  [Cab10b], which was derived in





**Figure 6.3:** Exclusivity graph  $G$  of KCBS scenario. If  $P(01 | i, i+1) = 1$ , then  $P(01 | i-1, i) = 0$  and  $P(01 | i+1, i+2) = 0$ , with sum taken modulo 5. In this sense, events connected by dashed lines are exclusive. Red dots indicate a maximal independent set of  $G$  (see main text).

1979 without any reference to QM:

$$p_{\text{QM}} \leq \vartheta(G)$$

For a pentagon graph,  $\vartheta(G) = \sqrt{5}$  and so in KCBS scenario, QM saturates the Lovasz bound.


Using this framework, we can make interesting comparisons to Bell inequalities. Figure 6.4 (left) shows the compatibility structure of CHSH scenario, which is also a cycle, but with only  $N = 4$  observables. We can construct the exclusivity graph by connecting events where a different outcome ( $a$  or  $b$ ) is obtained for the same measurement ( $x$  or  $y$  respectively). For example, the first term of the Bell witness in Eq. (6.21) is  $P(00 | 00)$  is exclusive with three other terms in the witness:  $P(11 | 00)$ ,  $P(11 | 01)$  and  $P(11 | 10)$ . The full exclusivity graph for Bell witness  $4p$  is shown in Figure 6.4 (right). Red dots highlight a maximum independent set of  $G_{\text{CHSH}}$ , which has size:

$$\alpha(G_{\text{CHSH}}) = 3$$

As a result, the deterministic NC bound for CHSH is given by:

$$4p_{\text{NC}} = 3 \implies p_{\text{NC}} = \frac{3}{4}$$

## 6. QUANTUM CONTEXTUALITY



**Figure 6.4:** (left) Compatibility structure of CHSH scenario. Alice’s and Bob’s measurements connected by solid lines are compatible due to NS assumption. (right) Exclusivity structure of CHSH scenarios. Events connected by dashed lines are exclusive.

in agreement with the deterministic locality bound in Eq. (6.30). So indeed, as argued before, NC extends the assumption of locality onto local systems. A second interesting observation is that for  $G_{\text{CHSH}}$  the Lovasz number

$$\vartheta(G_{\text{CHSH}}) = \frac{8}{\sqrt{5}}$$

and so the Lovasz bound on  $p_{\text{QM}}$  equals


$$p_{\text{QM}} \leq \frac{2}{\sqrt{5}} \approx 0.894$$

which is more than the Tsirelson’s bound (Eq. (6.33)) of  $p_{\text{QM}} \approx 0.85$ . In other words, CHSH scenario does not saturate the Lovasz bound available to its exclusivity graph [Cab10b]. Within the exclusivity framework, KCBS is the most fundamental NC inequality, since it has the smallest exclusivity graph, and results in the largest gap between QM and NC predictions [Cab13]. Furthermore, all other scenarios exhibit a gap between QM and NC predictions if and only if their exclusivity graphs contain odd- $N$  cycles, with  $N \leq 5$ .

### 6.3.4. Generalised KCBS inequalities and non-contextual content

A common feature of KCBS and CHSH inequalities is that their compatibility structures form a simple cycle of  $N$  vertices, with  $N = 4$  for CHSH and  $N = 5$  for KCBS. It turns out that every simple cycle of compatibility (for  $N \geq 4$  vertices) corresponds to a NC inequality. This family of results is known as  *$N$ -cycle inequalities* [Ara13]. Chained Bell inequalities (Section 6.2.4) are members of this family for even  $N$ . In this section, we extend the KCBS scenario to  $N > 5$  observables. This allows us to construct all odd  $N$ -cycle inequalities, thus “filling in the gaps” left by chained Bell inequalities.

For odd  $N$ , it is very straightforward to extend the KCBS proof onto  $N > 5$ . Consider  $N$



**Figure 6.5:** Compatibility structures for  $N$ -cycle inequalities for  $N = 4, 5, 6, 7$ .  $N = 4$  corresponds to CHSH inequality.  $N = 5$  corresponds to KCBS inequality.  $N = 6$  corresponds to a chained Bell inequality where Alice and Bob can measure 3 observables each.  $N = 7$  corresponds to a generalised KCBS inequality.

measurements  $A_1, A_2, \dots, A_N$  with cyclic compatibility structure ( $A_i$  is compatible with  $A_{i\pm 1}$ , with sum taken modulo  $N$ ), as in Figure 6.5. Consider a witness

$$S_N = \sum_{i=1}^N \langle A_i A_{i+1} \rangle. \quad (6.55)$$

Every one of those odd- $N$  cyclic compatibility structures produces an exclusivity graph which also is a simple  $N$ -cycle. For an odd  $N$ -cycle graph  $G_N$ , the independence number

$$\alpha(G_N) = \frac{N-1}{2}.$$

Following the proof in Sec. 6.3.3:

$$S_N \geq N - 4\alpha(G_N)$$

which gives

$$S_N \geq S_{N,NC} = 2 - N$$

The quantum maximum can be constructed by adapting the  $N = 5$  construction as shown in Fig. 6.6. For any odd  $N$ , we will define measurement directions as

$$|\psi_i\rangle = U_i|0\rangle \quad \text{where} \quad U_i = R_z^{(i-1)(N-1)/2} \left( \frac{2\pi}{N} \right) R_y(\theta).$$

Pairwise orthogonality  $\langle \psi_i | \psi_{i+1} \rangle = 0$  occurs when

$$\theta = \theta_N = \arccos \sqrt{\frac{\cos(\pi/N)}{1 + \cos(\pi/N)}}.$$

## 6. QUANTUM CONTEXTUALITY



**Figure 6.6:** Measurement directions for a generalised KCBS scenario for  $N = 9$  (left) and  $N = 11$  (right) (c.f. Fig. 6.2). Orthogonal directions are connected with dashed lines.

Eq. (6.51) becomes

$$\langle A_i A_{i+1} \rangle = 1 - 4 \cos^2 \theta_N.$$

which gives

$$S_N = \sum_{i=1}^N \langle A_i A_{i+1} \rangle = S_{\text{QM}} = \frac{N - 3N \cos(\pi/N)}{1 + \cos(\pi/N)}.$$

The general limit for compatible measurements is the algebraic limit  $S_N \geq -N$ . This can be seen by simply noting that a model where for all  $i$

$$\langle A_i \rangle = \frac{1}{2} \quad \text{and} \quad \langle A_i A_{i+1} \rangle = -1$$

has pairwise compatibility. Thus we can summarise the results for odd  $N \geq 5$  as

$$\begin{aligned} S_N &\geq S_{\text{N,NC}} = 2 - N && \text{for deterministic NC models} \\ S_N &\geq S_{\text{N,QM}} = \frac{N - 3N \cos(\pi/N)}{1 + \cos(\pi/N)} && \text{for compatible measurements in QM} \\ S_N &\geq S_{\text{NS}} = -N && \text{for compatible measurements} \\ S_N &\geq -N && \text{for incompatible measurements.} \end{aligned}$$

These results have some striking similarities for chained Bell inequalities in correlator form discussed in 6.2.5. The witness  $E_N$  in Eq. (6.41) has the same form as  $S_N$  in Eq. (6.55), except one of the correlators has a factor of  $(-1)$ . The bounds  $E_L$  and  $E_{\text{NS}}$  also have the same form, up to a factor of  $(-2)$ . The factor of 2 occurs since, in the Bell scenario,  $N$  referred to the number of measurements available to Alice and Bob separately<sup>8</sup>, meaning

<sup>8</sup> I used this notation since it is the most common

that the total number of different measurement directions is  $2N$ . In other words, a chained Bell scenario where Alice and Bob can each perform one of  $N$  measurements corresponds to a compatibility cycle of size  $2N$ .

To make the connection formal, for  $N \geq 4$ , we define the  $N$ -cycle witness

$$\Omega_N = \begin{cases} E_{N/2} & \text{for even } N \\ -S_N & \text{for odd } N \end{cases}$$

Since NC generalises the concept of locality, and compatibility generalises the idea of NS, we can write the combined bounds as

$$\begin{aligned} \Omega_N &\leq \Omega_{N,\text{NC}} = N - 2 && \text{for deterministic NC models} \\ \Omega_N &\leq \Omega_{N,\text{QM}} && \text{for compatible measurements in QM} \\ \Omega_N &\leq \Omega_{N,\text{NS}} = N && \text{for compatible measurements} \end{aligned}$$

where

$$\Omega_{N,\text{QM}} = \begin{cases} N \cos\left(\frac{\pi}{N}\right) & \text{for even } N \\ \frac{3N \cos(\pi/N) - N}{1 + \cos(\pi/N)} & \text{for odd } N \end{cases}$$

We generalise nonlocal content as *contextual content*

$$(1 - q)\Omega_{N,\text{NC}} + q\Omega_{N,\text{NS}} = \Omega_{N,\text{QM}} \quad (6.58)$$


which gives

$$q = \begin{cases} 1 - N \sin^2\left(\frac{\pi}{2N}\right) & \text{for even } N \\ 1 - N \tan^2\left(\frac{\pi}{2N}\right) & \text{for odd } N \end{cases}$$

These expressions result in extremely similar values of  $q$  for odd and even  $N$ , as plotted in Figure. 6.7. For both odd and even  $N$ ,  $q \rightarrow 1$  as  $N \rightarrow \infty$ . However, since for  $x \in (0, \pi/2)$  it holds that  $\tan(x) > \sin(x)$ , the curve of contextual content for odd- $N$  inequalities is always slightly below the one for even  $N$ .

## 6.4. State-independent contextuality

All NC inequalities presented so far are *state-dependent*. This means that the value of the witness, and hence the amount of violation, depends on the measured state. For example, the Bell witness in Eq. (6.21) reaches the Tsirelson's bound (Eq. (6.33)) for a Bell state  $|\Phi^+\rangle$ , while the KCBS witness in Eq. (6.46) saturates the quantum limit when the measured state is  $|0\rangle$ . However, the reason behind the state dependence is different. In Bell scenarios,



**Figure 6.7:** Contextual content  $q$  vs number of measurements settings  $N$ . The chained-Bell scenario (even  $N$ , red points) results in values of  $q$  very similar to those in the generalised KCBS scenario (odd  $N$ , blue points).

where entanglement is a resource, maximally entangled states provide more value than partially entangled states, and non-entangled states are deemed classical. However, in local scenarios, it makes little sense to think of  $|0\rangle$  as more or less classical than  $|1\rangle$ . The reason measurements of  $|0\rangle$  violate the KCBS inequality, while measurements of  $|1\rangle$  do not, is the choice of measurement directions in Eq. (6.49). That specific selection was arbitrary, and in fact any state  $|v\rangle = V|0\rangle$  can maximally violate a KCBS inequality for measurement directions  $|\psi_i\rangle = U_i V|0\rangle$ . In Bell scenarios, non-classicality is due to entanglement, while in local contextual scenarios, non-classicality is due to measurements. This point is made explicit by *state-independent NC* inequalities.

#### 6.4.1. Yu-Oh construction

The smallest possible set of measurement directions which yields state-independent contextuality was found by [Yu12]. It considers a set  $V$  of 13 measurement directions  $|v\rangle$  on the qutrit sphere, with measurement outcomes  $A_v \in (-1, 1)$ . Following the original paper, we define the observable for the measurement along  $|v\rangle$  as:

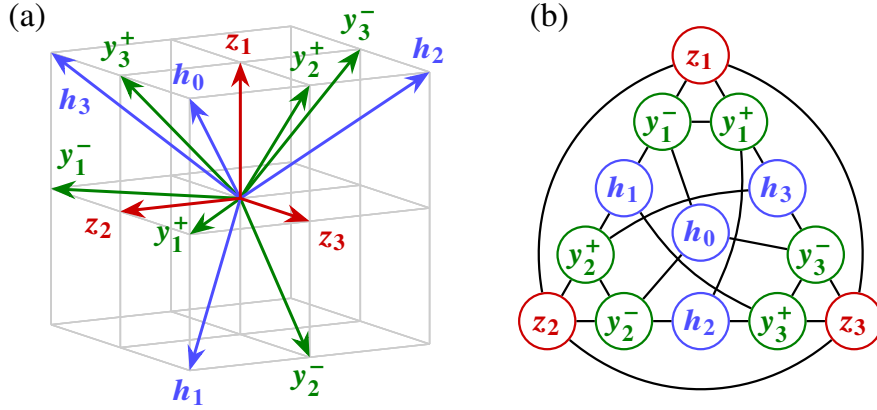
$$A_v = \mathbb{1} - 2|v\rangle\langle v|,$$

i.e. projection along  $|v\rangle$  corresponds to the outcome  $(-1)$ , opposite to the standard notation (Eq. (6.13)) elsewhere in the thesis, and we use the same symbol for the measurement operator and its outcome. The measurement directions are shown in Figure 6.8 a) and listed explicitly in Tab. 6.3. While the construction seems significantly more confusing than the KCBS scheme, each measurement direction is orthogonal to three or four other directions. This can be seen in the compatibility structure shown in Figure 6.8 b).

Like KCBS, the Yu-Oh witness considers correlators of compatible measurements. The

**Table 6.3:** Yu-Oh measurement directions specified on the qutrit sphere  $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle + \gamma|2\rangle$ 

$v$	$(\alpha, \beta, \gamma)$	$v$	$(\alpha, \beta, \gamma)$
$y_1^-$	$(0, 1, -1)$	$h_1$	$(-1, 1, 1)$
$y_2^-$	$(-1, 0, 1)$	$h_2$	$(1, -1, 1)$
$y_3^-$	$(1, -1, 0)$	$h_3$	$(1, 1, -1)$
$y_1^+$	$(0, 1, 1)$	$h_0$	$(1, 1, 1)$
$y_2^+$	$(1, 0, 1)$	$z_1$	$(1, 0, 0)$
$y_3^+$	$(1, 1, 0)$	$z_2$	$(0, 1, 0)$
		$z_3$	$(0, 0, 1)$



**Figure 6.8:** (a) Yu-Oh measurement directions  $V$ . Note that all unit vectors were extended from the unit sphere onto the unit cube in order to make their orthogonality more explicit. (b) Compatibility structure of the Yu-Oh measurement directions. Every measurement has either three ( $H$ ) or four ( $Y$  and  $Z$ ) compatible contexts. There are 48 edges in total. Figure from [Leu17] ©2017 American Physical Society

calculations are simplified by defining an *adjacency matrix*

$$\Gamma_{uv} = \begin{cases} 1 & \text{if } \langle u | v \rangle = 0 \\ 0 & \text{otherwise.} \end{cases}$$

Then the sum

$$\sum_{u,v \in V} \Gamma_{uv} \langle A_u A_v \rangle$$

will add up correlators of compatible measurements, and ignore all others. The Yu-Oh witness is given by

$$\chi_{\text{YO}} = \sum_{v \in V} \langle A_v \rangle - \frac{1}{4} \sum_{u,v \in V} \Gamma_{uv} \langle A_u A_v \rangle. \quad (6.59)$$

## 6. QUANTUM CONTEXTUALITY

and is state-independent, i.e. evaluates to the same result for every input state. To prove this result, notice first that the set of measurement directions  $V$  can be divided into sub-sets  $h_i \in H, z_i \in Z, y_i^\pm \in Y$  which individually resolve the identity:

$$\sum_{h \in H} |h\rangle\langle h| = \frac{4}{3} \times \mathbb{1} \quad \text{and} \quad \sum_{y \in Y} |y\rangle\langle y| = 2 \times \mathbb{1} \quad \text{and} \quad \sum_{z \in Z} |z\rangle\langle z| = \mathbb{1}$$

Given that  $A_v = \mathbb{1} - 2|v\rangle\langle v|$ , this implies that

$$\sum_{h \in H} \langle A_h \rangle = \frac{4}{3} \quad \text{and} \quad \sum_{h \in H} \langle A_y \rangle = 2 \quad \text{and} \quad \sum_{h \in H} \langle A_z \rangle = 1 \quad (6.60)$$

regardless of which state is being measured. To handle the correlator in Equation (6.59), we can use Equation (6.20a)

$$\sum_{u,v \in V} \Gamma_{uv} \langle A_u A_v \rangle = \sum_{u,v \in V} \Gamma_{uv} (\langle A_u \rangle + \langle A_v \rangle - 1). \quad (6.61)$$

Examining the compatibility graph in Figure 6.8 b) we see that measurement directions  $z_i$  and  $y_i^\pm$  all have four neighbours, while  $h_i$  have three. Therefore

$$\sum_{u,v \in V} \Gamma_{uv} (\langle A_u \rangle + \langle A_v \rangle) = 8 \sum_{y \in Y} \langle A_y \rangle + 8 \sum_{z \in Z} \langle A_z \rangle + 6 \sum_{h \in H} \langle A_h \rangle \quad (6.62)$$

In total there are 48 edges in the compatibility graph, so

$$\sum_{u,v \in V} \Gamma_{uv} = 48. \quad (6.63)$$

Putting Eqs. (6.61), (6.62) and (6.63) into Eq. (6.59) gives

$$\chi_{YO} = 12 - \sum_{y \in Y} \langle A_y \rangle - \sum_{z \in Z} \langle A_z \rangle - \frac{1}{2} \sum_{h \in H} \langle A_h \rangle$$

and using Eq. (6.60) implies that

$$\chi_{YO} = 8 + \frac{1}{3} \quad \text{in QM}$$

regardless of which state is being measured. To find a bound for deterministic NC models, it suffices to consider all  $2^{13}$  possible assignments of measurement outcomes  $\vec{a} = (A_{z_1}, A_{z_2}, \dots, A_{h_4}) = (\pm 1, \pm 1, \dots, \pm 1)$  (it is possible to construct an analytical proof,



but it is not particularly illuminating). Evaluating

$$\chi_{\text{YO}} = \sum_{i=1}^{13} a_i - \frac{1}{4} \sum_{i=1}^{13} \sum_{j=1}^{13} \Gamma_{ij} a_i a_j$$

for all possible vectors  $\vec{a}$  we find that

$$\chi_{\text{YO}} \leq 8 \quad \text{for deterministic NC models}$$

with equality achieved for 28 different assignments.

### 6.4.2. Optimal Yu-Oh inequality

The beauty of the Yu-Oh construction was the observation that a state-independent contextuality scenario can be constructed whenever a number of measurement averages sums to identity. In fact, the Yu-Oh observables allow one to specify many possible NC inequalities [Kle12]. Formally, we can define a general witness as:

$$\chi(\vec{\lambda}) = \sum_c \lambda_c \langle \Pi_{k \in c} A_k \rangle$$

where  $c$  specifies a set of mutually compatible directions, and  $\vec{\lambda} = (\lambda_{c1}, \lambda_{c2}, \dots)$  is a weight vector. It is possible to list all state-independent inequalities for the Yu-Oh rays by solving for

$$\chi(\vec{\lambda}) = 1. \tag{6.64}$$

For each solution, the NC deterministic bound:

$$\eta(\vec{\lambda}) = \sum_c \lambda_c (\Pi_{k \in c} a_k)$$

can be efficiently found by testing all possible assignments of  $\vec{a}$ .

The optimal Yu-Oh inequality can be constructed by solving for Eq. (6.64) while minimising  $\eta(\vec{\lambda})$ . This constitutes a linear program, and can be efficiently optimised. The final degree of quantum violation is quantified by

$$\mathcal{V} = \frac{\chi(\vec{\lambda})}{\eta(\vec{\lambda})} - 1.$$

For the original Yu-Oh witness,  $\chi = 8 + \frac{1}{3}$  and  $\eta = 8$ , resulting in  $\mathcal{V} = 1/24 \approx 4.2\%$

This optimisation leads to two interesting findings [Kle12]. First, the original Yu-Oh witness is not optimal among witnesses composed out of one- and two-point correlators. A different set of weights, listed explicitly in [Kle12], results in a doubled quantum violation of

## 6. QUANTUM CONTEXTUALITY

$\mathcal{V} = 1/12 \approx 8.4\%$ . Second, the bound can be further improved by using size-three contexts, i.e. sets of three mutually commuting measurement directions. Examining Fig. 6.8, we find four such sets:

$$\begin{aligned} C_3^A &= \{\{z_k, y_k^+, y_k^-\}\} \text{ for } k = 1, 2, 3 \\ C_3^B &= \{z_1, z_2, z_3\}. \end{aligned}$$

By allowing size-three contexts, we find an optimal inequality with a quantum violation of  $\mathcal{V} = 8/75 \approx 10.7\%$ . The optimal witness is explicitly given by:

$$\begin{aligned} \chi_{\text{opt}} &= \sum_{v \in H} 2 \langle A_v \rangle + \sum_{v \in V \setminus H} \langle A_v \rangle \\ &\quad - \sum_{(u,v) \in V \setminus C_2} 2 \Gamma_{u,v} \langle A_u A_v \rangle - \sum_{(u,v) \in C_2} \langle A_u A_v \rangle \\ &\quad - \sum_{(u,v,w) \in C_3^A} 3 \langle A_u A_v A_w \rangle, \end{aligned} \tag{6.65}$$

where  $C_2 = \{\{z_k, y_k^+\}, \{z_k, y_k^-\}, \{y_k^+, y_k^-\}\}$ . The violation evaluates to:

$$\begin{aligned} \chi_{\text{opt}} &= \frac{83}{3} \approx 27.667 && \text{in QM} \\ \chi_{\text{opt}} &\leq 25 && \text{for deterministic NC models.} \end{aligned}$$

The optimal inequality makes it easier to experimentally resolve the difference between the quantum and classical predictions, as we will see in Chapter 7.

### 6.5. Compatibility loophole

As is well known, strict testing of Bell inequalities is made difficult by numerous loopholes [Bru14]. Loopholes allow one to explain experimentally observed correlations - which seem to violate local determinism - in terms of purely classical models. These models may not be common sense, but as long as they can be applied, locality cannot be disproven beyond doubt. It took over 40 years from the first violation of Bell inequalities [Fre72] to the first loophole-free violation of Bell inequalities [Hen15]. This was made possible by both experimental developments - such as making more efficient single-photon detectors - and theoretical progress - such as finding out how to close the detection loophole without perfectly efficient detectors.

Quantum contextuality tests inherit certain loopholes from the Bell tests. For example, the argument about a detection loophole can be applied identically in both scenarios and overcome with high detection efficiency [Bru14]. However, the story of experimental loopholes in contextuality tests is overall very different. This is because of the *finite-precision loophole*, also known as the *compatibility loophole*, with the potential to invalidate

any experimental contextuality test.

### 6.5.1. The finite precision debate

The debate on whether local contextuality can be experimentally tested at all began at the turn of the millennium. In 1999, David A. Meyer published an article entitled “Finite Precision Measurement Nullifies the Kochen-Specker Theorem” [Mey99]. This article, together with [Ken99] and [Cli00], brought a powerful argument against experimental attempts to violate NC inequalities. The essence of the argument is the following. Every experiment has finite precision. Therefore, every time a projective measurement  $A$  is performed, producing a result with precision  $\epsilon$ , there is another projective measurement  $B$  which satisfies

$$||A - B|| < \epsilon$$

and so we cannot experimentally prove that we measured  $A$  rather than  $B$ . Following this observation, Meyer, Kent and Clifton demonstrated that, for every set of observed correlations and every  $\epsilon$ , there exists a deterministic NC model which reproduces the observed outcomes up to  $\epsilon$ . Therefore, as long as experiments have a finite precision, they cannot disprove deterministic NC models.

The reason this loophole does not apply to non-local Bell tests is two-fold. First of all, the nonlocality assumption does not rely on measurement calibration, only space-like separation. Secondly, the measurements in Bell tests need not be projective. In fact, while noisy measurements are naturally non-deterministic, the extended Bell’s theorem (Sec. 6.2.2) accommodates non-deterministic scenarios. There is no analogue of the extended Bell’s theorem in local scenarios, which can invalidate contextuality tests in the presence of experimentally unavoidable noise.

This “finite precision loophole”, nowadays usually referred to as the “compatibility loophole”, sparked a wave of rebuttals, with paper titles such as “Nullification of the Nullification” [App01], “Finite-precision measurement does not nullify the Kochen-Specker theorem” [Cab02] or a tongue-in-cheek “Finite precision measurement nullifies Euclid’s postulates” [Per03]. The rebuttals were themselves rebutted in [Bar04], and the rebuttals were counter-rebutted in [Cab09; Cab10a].

### 6.5.2. Experimental implications

By today, everyone seems to have formed an opinion about this problem, but no common agreement was reached. Instead, a number of parallel approaches have been developed in the literature.

### Exclusivity framework

The exclusivity framework (described in Sec. 6.3.3), also called the “Cabello-Severini-Winter framework” or the “Graph-theoretic framework”, claims that the traditional notion of contextuality can be an experimentally testable theory [Win14; Kun19].

The main assumption of this framework is that all measurements are fundamentally *sharp*. This assumption, defined in [Chi14b], means that the measurements are minimally disturbing and their outcomes are the same if performed repeatedly. In quantum theory, sharp measurements are represented by projectors. Following this assumption, the compatibility loophole can be addressed by assuming specific noise sources. There is no work addressing the loophole in a general local scenario, though it was shown theoretically [Cab10a] and experimentally [Hu16] that the loophole can be avoided in a nonlocal two-qutrit scenario.

Furthermore, the exclusivity framework requires the measurements to be performed sequentially, rather than simultaneously [Cab09]. The reason is that, in order for NC assumption to make sense, the measurement sequence itself needs to be context-independent. If that is not the case, the measurements are trivially contextual, and the NC assumption is baseless [Bar04]. A final advantage of the sequential measurement method is that it allows for extensive testing of hypotheses of measurement compatibility and measurement sharpness. We can check if  $A_i$  and  $A_j$  are compatible by measuring  $A_i$  alone, before  $A_j$ , after  $A_j$  or interleaved in a sequence of  $A_i - A_j - A_i$ , etc. Measurement sharpness can be verified by measuring  $A_i$  multiple times in a row and correlating the outcomes.

### Contextuality-by-default framework

This framework, developed predominantly by J.A. Larsson, E. N. Dzhafarov and J. V. Kujala [Lar02; Dzh12; Kuj15], explicitly avoids the assumption of perfect measurement compatibility. Instead, the experimentally measured amount of signalling is used to bound the amount of incompatibility of the underlying model. This is a powerful framework for analysing experimental data, and we apply it in Chapter 7 to analyse the KCBS experiment. We do not apply it to the Yu-Oh data, since no analysis to date reported on how it can be used in a state-independent scenario.

My personal difficulty with the contextuality-by-default framework is its philosophical opaqueness. The publications discussing its logical underpinnings are long and convoluted, and I never managed to understand them well enough to describe the method in my own words<sup>9</sup>. I have been told that the measurement sharpness assumption is still necessary to apply this framework, but I have been unable to find a conclusive answer one way or another. The same confusion remains about the issue of sequential versus simultaneous measurements.

---

<sup>9</sup> And that’s coming from someone who used to read 20th century German philosophy for fun!

### Spekkens’s operational framework

An explicitly operational framework of contextuality was pioneered by R. Spekkens in [Spe05] and later developed together with R. Kunjwal [Kun15; Kun19]. Within this approach, experimental states are not assumed to match any specific mathematical form. Instead, states are only described in terms of the actions performed to prepare them. The framework then allows for experimental noise and imperfections using the formalism of general probabilistic theories.

The operational approach is very methodologically sound and experimentally applicable [Maz15], although the operational KCBS inequality was only derived after the completion of our experiments [Kun19]. However, it is worth pointing out that the meaning of contextuality in the operational sense is very different from the traditional meaning, and it is difficult to say exactly how they relate. For example, while (traditional) contextuality requires at least 3 dimensions (qutrit), Spekkens’s contextuality occurs already for a single qubit. Likewise, it is unclear whether the computational implications of contextuality extend to the operational one. Finally, while traditional contextuality experiments aim at disproving NC deterministic theories, operational contextuality experiments disprove certain classes of generalised probabilistic theories. For those reasons, we have not used the operational approach to analyse our experiments.

### Sheaf-theoretic framework

The sheaf-theoretic framework, pioneered by S. Abramsky and A. Brandenburger in [Abr11], uses methods of logic and information theory to define contextuality through a lens of consistency. Specifically, contextuality arises whenever the measurements are locally consistent but globally inconsistent. Within this framework, it is possible to quantify the degree of contextuality through a concept of the “contextual fraction” and relate it to computational advantage [Abr17]. The expression for the contextual fraction of an N-cycle scenario turns out to exactly match the contextual content  $q$  as defined in Eq. (6.58), and we use it in Chapter 7 to analyse the strength of quantum correlations.

However, the sheaf-theoretic framework does not currently address the compatibility loophole directly. Work towards that goal was ongoing as of 2018, but no result on this topic has been published since<sup>10</sup>.

## 6.6. Is contextuality silly?

The title of this section, borrowed from Mermin’s classic review on the topic [Mer93], hints at the questions usually left unanswered by experimental papers. NC inequality violation forces us to abandon non-contextual deterministic HV models. But is the NC assumption (Eq. (6.44)) really worth fighting for? This section is structured as follows. First, I present

---

<sup>10</sup> S. Abramsky, private communication

a toy model aimed at clarifying the distinction between contextual and non-contextual explanations. Then I argue – following Mermin – that the NC assumption is reasonable when applied to the quantum theory. Then I turn to the issues usually left unsaid. First, I describe the cases where the NC assumption fails, and give examples of how it can happen experimentally. I then use those examples to argue that QRNGs based on contextuality make little sense.

### 6.6.1. Toy model of contextual and non-contextual theories

First, let me present a toy scenario, outside of the realm of QM, to illustrate the difference between explanations that assume NC, and those that do not. Consider two students, Alice and Bob, who take many multiple-choice exams. Their answers are random variables with certain correlations. In order to determine if Alice and Bob cheat, we perform statistical tests to see if Alice performs better on questions that were previously given to Bob, and vice versa. However, we find no such correlations, indicating that – in the language of contextuality – their answers are compatible. Does it mean that they are not cheating?

- A NC model would say so – since their answers don't influence each other, Alice and Bob probably did not communicate.
- A contextual model would say, on the other hand, that Alice and Bob discuss questions with each other, but this communication is not observable to you. This could happen either because they conspire to stay undetected<sup>11</sup>, or accidentally due to their incompetence.

Both types of models are allowed by physics, so is there a reason to believe that QM is described by a NC model? I struggle with this question myself. I believe that, on one hand, there are excellent reasons to believe that the fundamental theories of nature should be NC. At the same time, it is absolutely clear that contextual behaviours are common in nature. If contextual correlations explain outcomes of incompatible measurements, they might also affect some compatible measurements. Let us study those reasons in detail.

### 6.6.2. When assuming Non-contextuality makes sense

The main argument in favour of NC is that, if some quantum measurements admit a contextual model, then this model will be non-local when the measurements become space-like separated [Mer93; Bud21]. Consider once again the Bell scenario. A violation of Bell inequality forces us to abandon either determinism or locality, and most physicists choose to abandon determinism. Now imagine Alice and Bob move closer together, such that while their measurements are still compatible, neither locality nor NS are necessary assumptions. It stands to reason that their correlations are still explained by outcome indeterminism, and that locality equations – now called NC equations – still hold. If we abandon NC, we either

---

<sup>11</sup> A common strategy in the game Werewolves

abandon locality when measurements become space-like separated, or we must claim that as soon as the measurements are conducted far enough apart, their nature changes from deterministic to indeterministic. This, in my view, is the strongest argument in favour of a NC theory of nature. Returning to our toy model, if Alice and Bob get the same results whether they solve the exams in the same room or on different planets, then locality gives us a strong reason to believe they did not communicate. While the KCSB and the Yu-Oh scenarios are necessarily local, since they involve measurements of individual multi-level systems, multi-particle scenarios allow for explicit mapping between contextual and Bell scenarios [Mer93; Kir09b].

Aesthetic and philosophical considerations may also speak in favour of NC. Abandoning NC would mean observed measurement compatibility is a result of averaging over multiple physical states in which those measurements are not compatible. It would certainly be an odd theory, and not be favoured by Occam's razor. However, we know many emergent phenomena in nature - randomness can arise from chaotic behaviour in (deterministic) classical mechanics, and deterministic classical mechanics can, in turn, arise from non-deterministic quantum phenomena<sup>12</sup>. So compatibility emerging out of non-compatibility would be weird, but plausible.

### 6.6.3. When assuming Non-contextuality makes little sense

While it is plausible that the fundamental theory explaining quantum phenomena obeys NC, it does not mean that everything does. The critical - and yet often neglected - difference between the assumptions of NC and locality is that faulty lab equipment, badly designed experiment or conspiring adversaries can easily break NC, but still obey locality. Let me give two illustrations of this problem, and analyse how they break the non-contextuality assumption.

#### Amplifier noise

Consider a situation in which different settings  $(x, y)$  correspond to control pulses of ideally the same amplitude  $A$  but different length  $(t_x, t_y)$ . The measurement outcomes  $(a, b)$  are therefore determined by the pulse area.

Suppose that a noise with a single frequency component  $\omega$  modifies the amplitude to  $A(1 + \epsilon \cos \omega t)$ . Every shot, the area of the first pulse depends on the phase  $\phi$  of the noise at the start of the pulse. We can therefore consider  $\phi$  to be a hidden variable for the problem:

$$P(a | xy) = \sum_{\phi} P(\phi) P(a | x\phi).$$

Assuming  $\phi$  is uniformly distributed, the average pulse area remains  $At_x$  for the first pulse.

Whether or not this scenario admits a non-contextual model depends on how the second

---

<sup>12</sup> Debatable: consciousness emerges in the brain from non-conscious behaviour of individual cells

## 6. QUANTUM CONTEXTUALITY

pulse is performed. When the second pulse is fired, the noise phase  $\phi' = \phi + \omega(t_x + t_{\text{delay}})$ . Suppose  $t_{\text{delay}}$  is kept constant. Then each shot the area of the second pulse depends on the length of the first pulse:

$$P(b | xy\phi) \neq P(b | y\phi).$$

However, this effect is not visible after averaging over uniformly distributed  $\phi$ :

$$P(b | xy) = \sum_{\phi} P(\phi) P(b | xy\phi) = P(b | y).$$

Therefore, the resulting pulses satisfy compatibility (Eq. (6.43)) but are contextual. In other words, the noise will create extra correlations indistinguishable from the quantum ones.

One way to combat this contextuality is to adjust  $t_{\text{delay}}$  depending on  $t_x$ , such that  $t_{\text{delay}} + t_x$  is kept constant. Then  $\phi'$  is independent of  $t_x$ , and so:

$$P(b | xy\phi) = P(b | y\phi).$$

In such cases, the non-contextually condition (Eq. (6.44)) is satisfied. However, removing one vulnerability creates another. For example, the temperature of the amplifier at the start of the pulse can influence the outcome. It is realistic to expect the temperature at the start of the second pulse to depend directly on  $t_{\text{delay}}$ . When  $t_{\text{delay}}$  holds information about  $t_x$ , then those thermal transients can behave contextually.

### Broken detector

The previous example showed that temporal noise correlations can lead to contextuality without incompatibility. However, even simpler noise can create contextual effects. Consider again a sequence of measurements with settings  $(x, y)$  and outcomes  $(a, b) \in \{-1, 1\}$ . Suppose that the qubit measurement device is faulty and returns results deterministically, based only on its internal state  $\lambda$ . For simplicity, assume that it has only two equally probable states: when  $\lambda = 1$ , the measurement returns  $(+1, -1)$ , and when  $\lambda = -1$ , the outcome is  $(-1, +1)$ . In this case, measurements are compatible, since  $P(b|xy) = 1/2$  is independent of  $x$ .

It seems that the system is also non-contextual since each shot the measured outcome depends only on  $\lambda$  and not on  $x$ . However, there is a caveat. Contextual inequalities involve sequential measurements of different pairs of compatible observables. We cannot guarantee that every observable is either always measured first or always measured second. Instead, the measurement order depends on the choice of compatible measurements.

This can be made explicit in the KCBS scenario. Consider the measurement setting  $y = x_i$ , which compatible with measurements  $x = x_{i-1}$  and  $x = x_{i+1}$  (Sec. 6.3.1). In general, all measurements need to occur in some order, and so whether  $x_i$  is measured first depends



on whether it is measured together with  $x_{i-1}$  or  $x_{i+1}$ . Suppose  $x_i$  is measured first when  $x = x_{i-1}$  and second when  $x = x_{i+1}$ . Then every shot the measurement of  $x_i$  gives

$$\begin{aligned} A_i &= -\lambda & \text{when } x &= x_{i-1} \\ A_i &= \lambda & \text{when } x &= x_{i+1} \end{aligned}$$

Therefore,  $P(b | xy\lambda) \neq P(b | y\lambda)$  and the system is contextual. As a result, this broken detector could easily violate the KCBS inequality and its extensions. Since it exhibits no incompatibility, it would also pass checks imposed by the contextuality-in-default approach [Kuj15]. Different measurements, such as sharpness tests, would be necessary to discover the fault.

### Experimental implications

These examples are not meant to imply that contextual noise is unavoidable, or even common. However, it is necessary to understand that compatible outcome will, in general, be caused by a mixture of contextual and non-contextual phenomena. As a result, it is never possible to fully exclude a contextual explanation of the measurement results.

Instead, the best an experiment can do is to diligently track its noise sources. Every experiment can be treated as an opportunity to disprove some contextual interpretation of the data. By performing the measurements with different experimental settings, we can find out whether the results agree with the QM calculation, or whether they point to an error source. Once a noise source is identified, we can study its properties by searching for contextual effects. Measurements can be conducted with their orders swapped to verify their projective nature, and measurement sharpness tests can be performed by repeating the same measurement many times over. All those methods are applied in the experimental tests reported in Chapter 7

The remaining weakness of these methods is the following: once all the noise sources are identified, there is no general method of including them in the analysis. This applies, for example, to the unavoidable detection errors (Sec. 5.3), which corresponds to a violation of the sharpness assumption. In Chapter 7 we quantify those errors for our experiments, and I hope that future theoretical developments will show how to include those effects in the analysis of the results.

#### 6.6.4. Why perform contextuality experiments?

Another weakness of contextuality experiments is the limited reach of their conclusions. Assuming NC to be a believable principle, and that all contextual noise sources are eliminated, a NC inequality violation forces us to conclude that the world is non-deterministic. Unfortunately, this is as far as the implications take us. There is currently no known local analogue to the extended Bell's inequality, which would force us to conclude that QM is in some sense non-contextual even when it's probabilistic [Spe14].

Given that non-determinism is already firmly established by the Bell inequalities, is there any good theoretical motivation to perform contextuality tests? In all honesty, it is difficult to answer this question with a resounding “yes”. While there are hints of interesting observations here and there [Win14], the clear reason to perform such experiments has yet to be spelt out.

The following informal goals motivated the experiments in the upcoming chapter:

- **KCBS experiment.** As it is the smallest local contextual scenario, and one with the largest quantum violation, we study the extent to which the quantum bound (Eq. (6.48)) can be experimentally saturated. The motivation for this study comes from the field of Bell inequalities, where it was suggested that Tsirensen’s bound may not be fully achievable in nature [Gri15] (this conjecture was later disproved experimentally [Poh15]). We hope that our experiment can guide future theoretical developments on whether the KCBS bound can be fully saturated.
- **Generalised KCBS experiment.** By measuring the contextual content  $q$  in odd-cycle scenarios for increasing  $N$ , we complement the work on chained Bell inequalities. We hope that future theoretical work will clarify whether the contextual content can constrain post-quantum theories in the same way as the non-local content does [Col11] and whether it has any direct information-theoretical meaning [Abr17].
- **Yu-Oh experiment.** We perform the Yu-Oh experiment in a state-recycling fashion [Waj16]. This experiment explores a different way of generating quantum correlations. Furthermore, by randomising the measurement directions of the fly, we address a different class of HV models than other experiments do.

### 6.6.5. Random number generation?

Random number generation has long been listed as a possible application for quantum contextuality [Cal09; Abb10; Abb12]. It has been stated that contextual inequalities provide some of the security of DI QRNGs based on Bell’s theorem [Pir10], without the overheads that make those non-local sources impractical [Her17; Kul17; Um20]. This motivated multiple experimental demonstrations of the phenomenon [Um13; Kul17; Um20]

These assertions are frankly baffling and do not stand up to scrutiny [Hor10]. I think it is more accurate to re-phrase them to say: *contextuality provides essentially none of the security of Bell’s theorem, but inherits some of its overheads.*

Contextuality-based randomness generators implicitly make the following observation: speaking precisely, NC inequality violations do not imply that all of QM is non-deterministic, but only that QM measurements give random outcomes during contextuality tests (assuming model non-contextuality). By continuously running a contextuality test, we can therefore verify that the outcomes are always genuinely random. However, the first question of any randomness analysis should be: “random to whom?” [Sca19]. The main argument against contextual QRNGs is that they are insecure against adversaries who, like Alice and Bob in

Sec. 6.6.1, are of course allowed to behave contextually. Furthermore, any contextual noise (Sec. 6.6.3) provides a loophole through which an adversary can predict the measurement results without access to the laboratory.

On the other hand, QM allows for the construction of simple and high-throughput QRNGs, for example, based on photon shot noise of coherent states [San14], or single photons impinging a beamsplitter [Jen00]. It is difficult to conjecture<sup>13</sup> any remotely realistic agent who can “hack” such a simple generator, but from whom a contextual QRNG remains secure. It is likewise difficult to propose a set of physical laws where randomness emerges only during contextuality tests.

Future demonstrations to contextual advantage to QRNG must seriously examine the underlying security assumptions. It is insufficient to claim protection from non-contextual adversaries since that is not a realistic behaviour constraint.

---

<sup>13</sup> Try it!



## 7. Experimental studies of quantum contextuality

This chapter reports the experimental results of quantum contextuality tests, as well as details of multi-level quantum operations. It reports and expands on the findings originally published in:

1. F. M. Leupold, M. Malinowski, C. Zhang, V. Negnevitsky, J. Alonso, J. P. Home and A. Cabello, “Sustained state-independent quantum contextual correlations from a single ion”, *Physical Review Letters* **120** (2017).
2. M. Malinowski, C. Zhang, F. M. Leupold, V. Negnevitsky, J. Alonso, J. P. Home and A. Cabello, “Probing the limits of correlations in an indivisible quantum system”, *Physical Review A* **98** (2017).

By and large, figures and text in this chapter is subject to the copyright of American Physical Society ©2017, and are included in accordance with the permitted terms of use [APS17]. Some of the figures and text included in this chapter were originally prepared for publication by F. M. Leupold.

### Experimental setup

The experiments described in this chapter were performed with a single  $^{40}\text{Ca}^+$  ion in the setup 1.0. While Chapter 5 only discussed operations in the traps with integrated waveguides in the setup 2.0, the single-ion performance in the setup 1.0 was comparable. The main differences relevant for this chapter are:

- Optical qubit rotations are performed at Rabi frequencies of  $\Omega \approx 2\pi \times 10$  kHz or less with free-space  $\lambda = 729$  nm light propagating at 45 deg to the quantization axis. The depth of Rabi oscillations is limited to  $\approx 1 - 2 \times 10^{-3}$ , likely due to phase- and amplitude-modulation from cryogenic vibrations.
- No current-carrying wire is available to address the Zeeman qubit directly.
- Optical qubit coherence decays at a timescale of  $T_{\phi_2} \approx 2.2$  ms (Gaussian), limited by the cryogenic vibrations.
- The axial centre-of-mass mode is at  $\omega = 2\pi \times 2$  MHz and EIT cooling prepares it in  $\bar{n} \approx 0.2$ . The mode exhibits heating rate in the range of  $\dot{n} = 200 - 1000$


quanta per second. This electric-field noise was likely technical in origin since it was independent of the trap chip temperature up to 200 K, but our attempts at reducing or stabilising it were unsuccessful. The radial mode temperatures and heating were never characterised.

- The photon collection efficiency was  $\approx 50\%$  lower than in the setup 2.0, resulting in increased detection errors (Sec. 7.1.3).

### 7.1. Ion as a qutrit

The NC inequalities rely on correlations between sequential qutrit measurements. In this section, we discuss the experimental ingredients necessary to record and projectively measure a qutrit encoded within a single trapped  $^{40}\text{Ca}^+$  ion.

Long-lived electronic sub-levels of  $^{40}\text{Ca}^+$  provide a near-perfect platform for encoding higher-dimensional quantum information. Within a single ion, it is possible to use up to  $d = 7$  levels at a time, six within the  $D_{5/2}$  Zeeman manifold, and one in the  $S_{1/2}$  sub-space. The remaining  $S_{1/2}$  sub-level can be used to mediate qutrit control, but both of the  $S_{1/2}$  states become mixed during a fluorescence readout, effectively reducing the useful dimensionality. In this section, we demonstrate the single-qutrit control through an example of a single qutrit ( $d = 3$ ), but the techniques described here generalise readily to  $d = 4, 5, 6, 7$ . A reader interested in the details and extension of those techniques is recommended to consult [Low19] which discusses the theory of universal single- and two-qutrit control in  $^{137}\text{Ba}^+$ .



**Figure 7.1:** Energy level diagram of the  $^{40}\text{Ca}^+$  ion. Qutrit states  $|0\rangle$ ,  $|1\rangle$ , and  $|2\rangle$  are encoded in Zeeman sub-levels highlighted in red. Coherent qutrit rotations between them are achieved with laser pulses at  $\lambda = 729\text{ nm}$ . Fluorescence measurements using an excitation laser at  $\lambda = 397\text{ nm}$  and  $\lambda = 866\text{ nm}$  (not shown) project the qutrit state into either  $|0\rangle$  (“bright”) or the  $|1\rangle, |2\rangle$ -manifold (“dark”).

The qutrit encoding used in this work is shown in Fig. 7.1. Quantum information is stored in the electronic levels  $|0\rangle = |S_{1/2}, m_j = -1/2\rangle$ ,  $|1\rangle = |D_{5/2}, m_j = -3/2\rangle$  and  $|2\rangle = |D_{5/2}, m_j = -1/2\rangle$  (note that we reversed the qubit notation from previous chapters, where  $|0\rangle$  was always the higher-energy state). Coherent single-qutrit rotations are implemented as sequences of  $\lambda = 729\text{ nm}$  pulses, tuned to resonance with either the  $|0\rangle \leftrightarrow |1\rangle$  or the  $|0\rangle \leftrightarrow |2\rangle$  transition, as described in Sec. 7.1.1 and Sec. 7.1.4. Standard optical qubit readout (Sec. 2.3.6) implements a binary measurement distinguishing  $|0\rangle$  from other qutrit states (Sec. 7.1.2). In Sec. 7.1.3 we describe how to make this measurement projective, allowing us to perform correlation measurements.

### 7.1.1. Coherent rotations

In full generality, the problem of single-qutrit control is the problem of generating an arbitrary  $d$ -dimensional unitary operation. Ignoring the global phase, the desired unitary is then a member of the  $d$ -dimensional special unitary group:

$$U \in SU(d).$$

and the available control Hamiltonians only couple two levels at a time. The simplest approach to synthesising  $U$  is through a sequence of resonant two-level rotations. To focus on the  $^{40}\text{Ca}^+$  level structure, suppose that the available couplings connect  $|0\rangle$  to all the other qutrit states. Extending the results from Sec. 2.3.1, the unitary of a resonant rotation on the  $|0\rangle \leftrightarrow |k\rangle$  transition is described by an operator:

$$R_k(\theta, \phi) = e^{-i\frac{\theta}{2}(\sigma_{k,x} \cos(\phi) + \sigma_{k,y} \sin(\phi))}, \quad (7.1)$$

where

$$\begin{aligned} \sigma_{k,x} &= |0\rangle\langle k| + |k\rangle\langle 0|, \\ \sigma_{k,y} &= -i|0\rangle\langle k| + i|k\rangle\langle 0| \end{aligned}$$

are the Pauli operators on the  $\{|0\rangle, |k\rangle\}$  subspace. The decomposition of  $U$  onto a sequence of  $R_k$  can be performed as described in [OLe06]. Formally, for any  $U$ , it is possible to use no more than  $d(d-1)/2$  rotations to algorithmically construct an operator  $U'$ :

$$U' = \prod_{i=1}^{d(d-1)/2} R_{k_i}(\theta_i, \phi_i).$$

which only differs from  $U$  by the phase of its diagonal entries. The final mapping from  $U'$  to  $U$  can be performed as a virtual phase update, as a sequence of  $(d-1)$  off-resonant phase-shifting pulses, or decomposed onto a sequence of additional  $2(d-1)$  rotations of the form in Eq. (7.1).

An interesting question is whether applying multiple drives in parallel can lead to a more

## 7. EXPERIMENTAL STUDIES OF QUANTUM CONTEXTUALITY

efficient decomposition. Indeed, theoretical work has shown that  $U$  can be decomposed into only  $d$  interaction steps (as opposed to  $\sim d^2$  above) if all transitions can be driven at once. This is achieved either through individually adjusting the Rabi frequencies and phases of all the drives for a given interaction time [Iva06], or by means of adiabatic passage [Rou13]. However, within this method,  $|0\rangle$  is only used as an ancilla and not as a computational state, reducing the available dimensionality. While interesting, this method is not beneficial for low values of  $d$ , and we have not pursued it experimentally.

We now focus on the qutrit case ( $d = 3$ ). Recall that in the KCBS and Yu-Oh contextuality tests, all the relevant measurement directions lie on the real-valued qutrit sphere. This makes the unitary synthesis straightforward and intuitive, as the problem selecting a qutrit observable reduces onto a problem of rotating unit vectors in the Cartesian space. Cartesian rotations can be generated by a combination of the following rotation matrices:

$$R_x(\theta) = \begin{pmatrix} \cos \theta & -\sin \theta & 0 \\ \sin \theta & \cos \theta & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad R_y(\theta) = \begin{pmatrix} \cos \theta & 0 & \sin \theta \\ 0 & 1 & 0 \\ -\sin \theta & 0 & \cos \theta \end{pmatrix}.$$

Relevantly, rotation by angle  $\theta$  around  $z$  is recomposed into

$$R_z(\theta) = R_x(-\pi/2)R_y(\theta)R_x(\pi/2) \quad (7.2)$$

Note that we swapped  $x \leftrightarrow z$  compared to traditional definitions to associate  $|0\rangle$  with the  $z$  axis and  $|1\rangle$  with the  $x$  axis, as in Sec. 6.3.1. Explicitly writing the rotation operators from Eq. (7.1) as single-qutrit matrices gives

$$R_1(\theta, \phi) = \begin{pmatrix} \cos\left(\frac{\theta}{2}\right) & -ie^{-i\phi} \sin\left(\frac{\theta}{2}\right) & 0 \\ -ie^{i\phi} \sin\left(\frac{\theta}{2}\right) & \cos\left(\frac{\theta}{2}\right) & 0 \\ 0 & 0 & 1 \end{pmatrix},$$

$$R_2(\theta, \phi) = \begin{pmatrix} \cos\left(\frac{\theta}{2}\right) & 0 & -ie^{-i\phi} \sin\left(\frac{\theta}{2}\right) \\ 0 & 1 & 0 \\ -ie^{i\phi} \sin\left(\frac{\theta}{2}\right) & 0 & \cos\left(\frac{\theta}{2}\right) \end{pmatrix}.$$

Therefore, we can implement Cartesian rotations through resonant pulses using the mapping:

$$R_x(\theta) = R_1(2\theta, \pi/2) = R_1(-2\theta, 3\pi/2),$$

$$R_y(\theta) = R_2(2\theta, 3\pi/2) = R_2(-2\theta, \pi/2).$$

where the middle term is used when  $\theta > 0$ , and the last term when  $\theta < 0$ . Of course, it is always possible to find a Cartesian rotation with  $\theta > 0$ , but this maps onto a longer pulse.

Writing a general measurement direction on the qutrit sphere as  $U|0\rangle = a|0\rangle + b|1\rangle + c|2\rangle$ ,



we find that  $U$  can be always written as a sequence of only two Cartesian rotations:

$$U = R_x(\alpha)R_y(\beta)$$

where the rotation angles are given by

$$\alpha = \pm \arccos\left(\frac{a}{\sqrt{a^2 + b^2}}\right), \quad \beta = \pm \arccos(c).$$


and the  $\pm$  coefficients depend on the signs on  $a$ ,  $b$  and  $c$ . Therefore, the measurement direction can be rotated to  $U|0\rangle$  using at most two resonant pulses. We will write the rotation sequence for a general unitary  $U$  as:

$$U = R_1(\theta_1, \phi_1)R_2(\theta_2, \phi_2). \quad (7.3)$$

In Sec. 7.2 and Sec. 7.4 we tabulate the explicit decompositions of specific measurement directions onto individual pulses. In Sec. 7.1.3 we discuss the remainder of the measurement sequence.

### 7.1.2. Multi-level readout

$D$ -level qudit readout can be performed as a sequence of  $d - 1$  fluorescence readout steps and  $d - 2$  shelving pulses. The first cycle of the sequence is shown in Fig. 7.2. The sequence



**Figure 7.2:** Qudit readout cycle. (A) Fluorescence measurement records the total population in  $S_{1/2}$  (B) Optical pumping moves the population from  $|a\rangle$  to  $|0\rangle$  (C) A shelving pulse moves the population from  $|1\rangle$  to  $|0\rangle$

begins with optical qubit readout (Sec. 2.3.6), which causes fluorescence with probability  $P(0)$  equal to the probability of finding the qudit in  $|0\rangle$ . The ground-state population is then re-initialised  $|0\rangle$ , emptying the other ground state  $|a\rangle = |S_{1/2}, m_J = +1/2\rangle$ . This is done through optical pumping on a dipole transition (Sec. 2.3.5), but with the  $\lambda = 854$  nm laser turned off. The population from the qudit state  $|1\rangle$  is then shelved into this ancillary ground state with a  $|a\rangle \leftrightarrow |1\rangle$  pi-pulse, and the measurement sequence is repeated to

obtain the combined population  $P(0) + P(1)$  of  $|0\rangle$  and  $|1\rangle$ . The sequence of optical pumping and shelving is repeated a total of  $d - 2$  times, with the final measurement yielding  $P(0) + P(1) + \dots + P(d - 2)$ . The population in  $|d - 1\rangle$  can be found by subtracting that number from 1.


While contextuality tests do not require direct multi-level readout, we used it frequently for error debugging. Furthermore, while only a single shelving pulse is needed for  $d = 3$ , measuring more levels is a handy tool for constraining undesired population loss outside of the qutrit subspace. However, the readout error increases with the number of shelving pulses, and measuring small populations is challenging due to errors in the optical pumping on a dipole transition, which leave  $\sim 1 \times 10^{-3}$  of the ground-state population in  $|a\rangle$ .

### 7.1.3. Projective binary readout

Projective binary readout combines the ideas from Sec. 7.1.1 and Sec. 7.1.2. The resulting measurement sequence along  $|\psi\rangle = U|0\rangle$  is shown in Fig. 7.3. The sequence begins with a unitary rotation  $U^\dagger$  which can be always written as

$$U^\dagger = (R_1(\theta_1, \phi_1)R_2(\theta_2, \phi_2))^\dagger = R_2(\theta_1, \phi_1 + \pi)R_1(\theta_2, \phi_2 + \pi).$$

The readout step consists of optical qubit readout and re-initialisation into  $|0\rangle$ . The readout is followed by a unitary  $U$  given by Eq. (7.3), which completes the sequence. Crucially,



**Figure 7.3:** Binary measurement of a qutrit

the re-initialisation step must be followed by ion re-cooling. This is because photon recoil from resonant fluorescence heats the ion’s motion, leading to errors on rotations following a bright readout [Neg18a]. We handle this challenge in two ways.

In the Yu-Oh experiment we perform the detection with a  $2\pi \times -10$  MHz detuning, corresponding closely to the optimal Doppler-cooling parameters (Sec. 5.6.2). We then adjust the detection time and amplitude to simultaneously collect enough photons and reduce subsequent pulse errors to acceptable levels.

Readout errors can be obtained by fitting the histogram of photon counts, as we did in Sec. 5.3.1. For this “cooling detection” find mean counts counts of (0.7, 18.8) in a detection window of 160  $\mu$ s. At a threshold of 5.5, the corresponding to readout errors are  $(2.0.1.9) \times 10^{-4}$  for the dark and bright states respectively (including the effect of dark state decay). However, this number only quantifies how well the measurement outcome matches the pre-measurement state, and not how well the post-measurement state matches these. To quantify that error, we perform experiments with the measurement  $\hat{M}^{[0]}$  repeated

twice in a row, measuring the conditional probability  $P(A_{|0\rangle} = m \mid A_{|0\rangle} = n)$  of the second measurement outcome  $m$  given the first measurement outcome  $n$ . With this method, we find:

$$\begin{aligned} P(A_{|0\rangle} = +1 \mid A_{|0\rangle} = -1) &\approx 5 \times 10^{-4} \\ P(A_{|0\rangle} = -1 \mid A_{|0\rangle} = +1) &\approx 5 \times 10^{-4}, \end{aligned}$$

which is approximately double the error of the histogram fit. Note that the decay errors affect the repeated measurement differently than a single measurement. For example, a decay event which causes an initially dark ion to appear bright in the first measurement leaves the ion bright for the second measurement, and thus does not appear as an error. On the other hand, an undetected decay event in the first measurement appears as an error of the repeated measurement, but not of a single measurement.

The rotation errors are characterised by measuring  $\hat{M}^{(1)}$  or  $\hat{M}^{(2)}$  following a measurement of  $\hat{M}^{(0)}$ . Ideally, the  $+1$  outcome of the first measurement implies a  $-1$  outcome of the second. Instead, our calibrations reveal

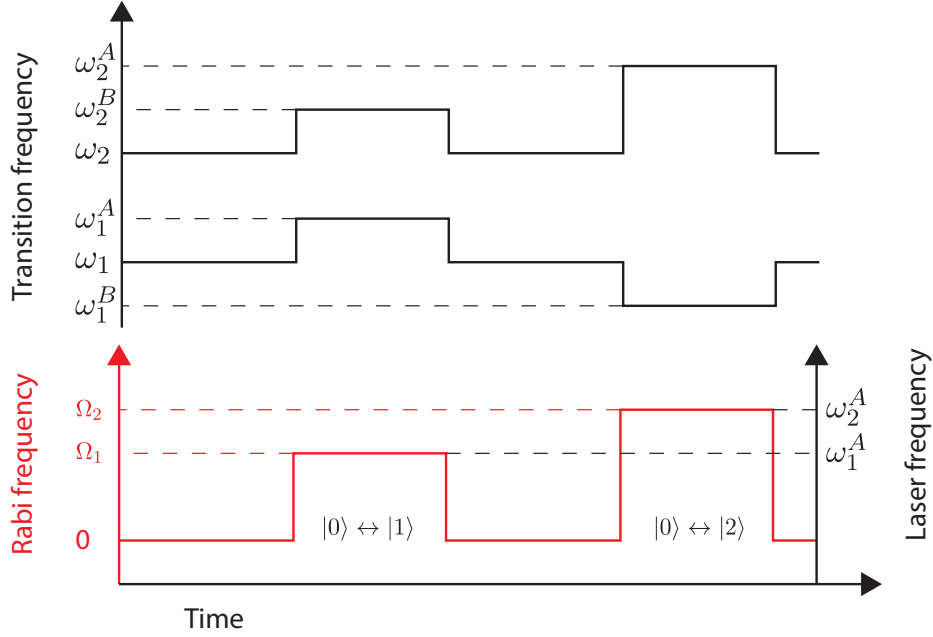
$$\begin{aligned} P(A_{|1\rangle} = +1 \mid A_{|0\rangle} = +1) &\approx 5 \times 10^{-3} \\ P(A_{|2\rangle} = +1 \mid A_{|0\rangle} = +1) &\approx 5 \times 10^{-3} \end{aligned} \tag{7.4}$$

consistent with a Doppler temperature of  $\bar{n} \approx 5$  on the axial mode (Sec. 5.5).

In the KCBS experiment, we improve the sequence two-fold. We perform the detection sequence with mean counts of  $(1, 25)$  in a window of  $200 \mu\text{s}$ , decreasing the dark/bright detection errors to  $(1.4, 0.2) \times 10^{-4}$  respectively at a threshold of 8.5. Each detection is still performed with settings close to Doppler cooling, but this time we follow each readout with a short ( $50 \mu\text{s}$ ) EIT cooling sequence. This cools the axial mode to  $\bar{n} \approx 0.2 - 0.3$  following a bright detection, close to the pre-measurement temperature. During a dark detection, the axial mode temperature increases by  $\approx 0.05 - 0.25$  quanta due to a heating rate of  $\dot{n} \approx 200 - 1000$  quanta per second. This reduces the pi-pulse errors (Eq. (7.4)) to the range of  $(1 - 2) \times 10^{-3}$ . This error is higher than expected given the measured temperature, and we conjectured it to be caused by Doppler shifts caused by cryogenic vibrations.

#### 7.1.4. AC Stark shift correction

Qutrit energy levels are affected by differential AC Stark shifts caused by the control pulses. This necessitates a phase correction when multiple rotations are applied in sequence. To that end, we keep track of three frequencies for each transition, as illustrated in Fig. 7.4. Denote the bare transition frequencies as  $\omega_1$  and  $\omega_2$ . When a resonant pulse is applied on the  $|0\rangle \leftrightarrow |1\rangle$  transition, these frequencies shift to  $\omega_1^A$  and  $\omega_2^B$  respectively. Similarly, a pulse resonant with the  $|0\rangle \leftrightarrow |2\rangle$  transition changes the transition frequencies to  $\omega_1^B$  and  $\omega_2^A$ . In general,  $\omega_i^A \neq \omega_i^B$ , both since the two resonant pulses have different powers, and



**Figure 7.4:** Illustration of the time evolution of qutrit transition frequencies. For each transition  $|0\rangle \leftrightarrow |i\rangle$ ,  $i = 1, 2$  the transition frequency shifts from  $\omega_i$  to  $\omega_i^A$  when driven resonantly, and to  $\omega_i^B$  when the other transition is driven resonantly.

due to differing quadrupole Stark shift contributions.

We calibrate the system as follows. First, we calibrate  $\omega_1^A$  and  $\omega_2^A$  by precise Rabi spectroscopy at the desired Rabi frequencies  $\Omega_1, \Omega_2$ , respectively. These frequencies are then used for the coherent rotations  $R_1$  and  $R_2$ . Then we calibrate  $\omega_1$  and  $\omega_2$  through Ramsey experiments (Sec. 5.4.2): a Ramsey delay of length  $\tau$  causes a phase shift of  $(\omega_i - \omega_i^A)\tau$  in the frame of the drive at  $\omega_i^A$ . Finally, we calibrate  $\omega_1^B$  by performing a Ramsey experiment on the  $|0\rangle \leftrightarrow |1\rangle$  transition, but with a  $|0\rangle \leftrightarrow |2\rangle$  rotation instead of a wait time. Specifically, a wait time of length  $\tau$  is replaced with a rotation

$$R_2\left(\frac{\Omega_2\tau}{2}, \pi\right) R_2\left(\frac{\Omega_2\tau}{2}, 0\right)$$

selected to always return the system to the  $\{|0\rangle, |1\rangle\}$  subspace. This rotation leads to a Ramsey phase shift of  $(\omega_1^B - \omega_1^A)\tau$  in the frame of the drive, and  $\omega_2^B$  is calibrated analogously.

In the Yu-Oh experiments, this correction was not taken into account due to overall higher pulse error (Eq. (7.4)), and the transitions were always assumed to be at their bare frequencies  $\omega_1, \omega_2$ . In the KCBS experiments, on the other hand, we applied the resonant pulses at the frequency  $\omega_i^A$  and added appropriate phase corrections proportional to the time spent with laser beams off or with the other rotation applied. However, during the data-taking stage, the Rabi frequencies were set low enough to keep all AC Stark shifts below  $2\pi \times 100$  Hz, making the corrections described in this section negligible.

The need for precise calibration of the bare frequency  $\omega_i$  is predominantly due to qutrit measurements. In a typical measurement sequence (such as Fig. 7.3), the majority of time ( $t_D = 170 \mu\text{s}$  in the Yu-Oh experiment and  $t_D = 260 \mu\text{s}$  in the KCBS experiment) is taken by the detection step (which includes readout, re-cooling and state preparation). The detection lasers only cause a common-mode AC Stark shift of the dark states, and so the relative energy difference between  $|1\rangle$  and  $|2\rangle$  is  $\hbar(\omega_1 - \omega_2)$ , the same as it is for free evolution. Thus, by finely calibrating the bare frequencies, we reduce phase errors associated with dark detections. Furthermore, every dark detection introduces an error due to finite coherence time of the  $|1\rangle \leftrightarrow |2\rangle$  transition. In setup 1.0, we measured the associated Gaussian decay time to be  $T_{\phi_2} \approx 10 \text{ ms}$ , resulting in an error of approximately

$$1 - e^{(t_D/T_{\phi_2})^2} \approx 6 \times 10^{-4}.$$

This error could be in principle reduced by modifying the detection sequence to include spin-echo pulses, although we did not attempt to do that as we remain dominated by coherent rotation errors.

## 7.2. Yu-Oh experiment


In this section, we discuss the experimental violation of the Yu-Oh inequality, which was derived and discussed in Sec. 6.4. This work followed a number of previous experimental violations of state-independent NC inequalities [Ams12; Mou10; Zu12; DAm13; Hua13; Cañ14a; Cañ14b]. In trapped-ion systems, previous works demonstrated contextual correlations in the Mermin-Peres square [Kir09b] and in the Yu-Oh scenario [Zha13]. All those demonstrations followed the same approach: i) prepare an input state, ii) measure multiple observables iii) repeat using all combinations of observables required for the test. State-independence was then argued by repeating the measurement of a number of different input states.

Our test follows a different principle, where the input state is not explicitly initialised, but effectively randomised by the measurement sequence itself. In this *state-recycling method* [Waj16], each measurement is performed on the state into which the system was projected by the previous measurement. This results in a continuous generation of quantum correlations, possible only in a state-independent scenario. In our experiment, we demonstrate it through a sequence of 53 million binary measurements, running autonomously over a period of 12 hours.

### 7.2.1. Measurement sequence

The measurement proceeds in a single sequence shown in Fig. 7.5. Each block measures an observable selected at random out of the 13 measurement directions  $v_i$  in the Yu-Oh construction as defined in Sec. 6.4. Each unitary  $U_{v_i}$  is decomposed onto a sequence of two

## 7. EXPERIMENTAL STUDIES OF QUANTUM CONTEXTUALITY



**Figure 7.5:** Sequential measurement scheme. A subsequence starts by initializing the ion state to  $|0\rangle$  and rotating it to the last ray from the previous subsequence,  $v_0$ . Every following ray measurement  $M_v$  then consists of a unitary transformation  $U_v^\dagger$  rotating the ray  $v$  onto  $|0\rangle$ , a projective measurement, and the back rotation  $U_v$ . The unitary transformations are written  $U_v = R_{1,v}R_{2,v}$ , where  $R_{2,v} = R_2(\theta_{2,v}, \phi_{2,v})$  and  $R_{1,v} = R_2(\theta_{1,v}, \phi_{1,v})$ . Subsequent measurement rays are determined by bit sequences from a QRNG, which are created after performing the respective previous projective measurement.

rotations as shown in Tab. 7.1.

**Table 7.1:** Pulse parameters for measurement directions  $v_i$  in the Yu-Oh set. The last column shows the corresponding bit sequence from the QRNG. If the QRNG delivers a bit sequence not present in this table, it is discarded and a new one is read in. Shorthand notation  $\theta_h = 2 \arctan(1/\sqrt{2})$  was used, and phases  $\phi_{i,v}$  were set to zero whenever  $\theta_{i,v} = 0$ .

$v$	$\theta_{1,v}$	$\phi_{1,v}$	$\theta_{2,v}$	$\phi_{2,v}$	QRNG
$y_1^-$	$\pi$	$\pi/2$	$\pi/2$	$3\pi/2$	0001
$y_2^-$	0	0	$3\pi/2$	$\pi/2$	0010
$y_3^-$	$\pi/2$	$3\pi/2$	0	0	0011
$y_1^+$	$\pi$	$\pi/2$	$\pi/2$	$\pi/2$	0100
$y_2^+$	0	0	$\pi/2$	$\pi/2$	0101
$y_3^+$	$\pi/2$	$\pi/2$	0	0	0110
$h_1$	$3\pi/2$	$\pi/2$	$\theta_h$	$\pi/2$	0111
$h_2$	$\pi/2$	$3\pi/2$	$\theta_h$	$\pi/2$	1000
$h_3$	$\pi/2$	$\pi/2$	$\theta_h$	$3\pi/2$	1001
$h_0$	$\pi/2$	$\pi/2$	$\theta_h$	$\pi/2$	1010
$z_1$	0	0	0	0	1011
$z_2$	$\pi$	$\pi/2$	0	0	1100
$z_3$	0	0	$\pi$	$\pi/2$	1101

The randomisation is done in real-time using a commercial QRNG<sup>1</sup>. It delivers a constant stream of random bits, from which we take groups of four and assign measurement directions  $v$  to them, as shown in Tab. 7.1. The random bits for an observable are created within a 50  $\mu$ s window after the detection event of the previous observable.

<sup>1</sup> Quantis from ID Quantique SA

### 7.2.2. Data acquisition and calibrations

Ideally, we would perform a single long series of measurements of randomly chosen observables. In practice, we interrupt the sequence to save collected data and periodically tune the control parameters. To sustain the sequence, we take subsequences containing a minimum of 1,000 measurements, which we interrupt when the last detection projected the qutrit onto  $|0\rangle$ . The next subsequence then starts by initializing the qutrit to  $|0\rangle$  and applying the rotation  $U_{v_0}$ , with  $v_0 = v_l$  the last ray from the previous sequence. In this way, all performed measurement sequences can be concatenated up to the 53 million in the final dataset. The periodic reset resulting from the repeated measurements with a finite probability to find the system in  $|0\rangle$  allows such long continuous sequences to be built up while restricting the propagation of calibration-related errors.

Each experimentally relevant parameter is automatically re-calibrated with a certain periodicity, and a script schedules a calibration experiment during the upcoming sequence interruption. The bare transition frequencies  $\omega_1$  and  $\omega_2$  are recalibrated every few minutes. Approximately every 10 minutes, we correct for laser intensity drifts by recalibrating all the laser rotation angles in Tab. 7.1. Every 30 minutes, we scan and calibrate the frequency of the  $\lambda = 397$  nm detection beam, and re-calibrate the histogram thresholds. The experiment runs continuously and autonomously until one of the lasers comes out of lock.

### 7.2.3. Purged data

Despite our best attempts at performing a continuous measurement, a certain amount of data cleaning was required prior to the analysis.

First of all, while preparing for the final data acquisition, we noticed an anomalous number of long sequences of dark measurements. We attributed this effect to off-resonant leakage into the states  $|D_{5/2}, m_j = -5/2\rangle$  and  $|D_{5/2}, m_j = +1/2\rangle$ , which are long-lived dark states outside our computational Hilbert space. This effect was very rare, hence difficult to study systematically. Therefore, we programmed the control system to detect these events in real-time and break, purging the subsequence and starting a new subsequence from the same  $v_0$  as was used for the purged subsequence. This reset step involves repumping the all the  $D_{5/2}$  states with a  $\lambda = 854$  nm laser into the  $S_{1/2}$  ground-state. We set the purge threshold to 55 dark measurements in a row, which should only occur with probability with  $50 \times 10^6 \times (2/3)^{55} = 1\%$  for 50 million measurements, which was our target dataset size. Indeed, during the data run, we recorded almost 200 such events within 53.5 million measurements.

The second data purge, which reduced the dataset size from 54.5 to 53.5 million measurements, was necessary due to an error in the live calculation of the ideal detection threshold during the experimental sequence. The computer control system consistently thresholded at 4.5 counts, whereas the ideal value was 5.5 counts (Sec. 7.1.3). Had we not split the single long sequence into subsequences of about 1000 measurements, this would not have been a problem as the identification of a “dark” or “bright” result can always be done

a posteriori based on the photon counts. However, 90 subsequences ended on a 5-count detection, from which a “bright” state was determined. We decided to omit those from the evaluation, as follows. If subsequence  $N$  ended on a 5-count detection and started with ray  $v_{0,N}$ , we ignored the corresponding line and all following ones until and excluding the next line that starts with  $v_{0,N}$ . The new  $N^{\text{th}}$  line can now again be correctly concatenated with the preceding one. While this reduced our dataset by around 1 million measurements, the impact on the numerical results was negligible.

### 7.3. Yu-Oh experiment results

#### 7.3.1. Data analysis

Every data point measured for an observable  $M_v$  consists of the measurement ray  $v$  and an outcome  $A_v = \pm 1$ . From the full data set, we collect the numbers  $N(A_v=a_1)$ ,  $N(A_u=a_1, A_v=a_2)$ , and  $N(A_u=a_1, A_v=a_2, A_w=a_3)$ , where  $A_u$ ,  $A_v$ , and  $A_w$  are successive measurements outcomes in that order, for all  $u, v, w \in V$  and all  $a_1, a_2, a_3 \in \{1, -1\}$ . Based on these numbers, we compute the expectation values

$$\begin{aligned}\langle A_v \rangle &= \frac{\sum_{a_1} a_1 N(A_v=a_1)}{\sum_{a_1} N(A_v=a_1)}, \\ \langle A_u A_v \rangle &= \frac{\sum'_{a_1, a_2} a_1 a_2 N(A_u=a_1, A_v=a_2)}{\sum'_{a_1, a_2} N(A_u=a_1, A_v=a_2)}, \\ \langle A_u A_v A_w \rangle &= \frac{\sum'_{a_1, a_2, a_3} a_1 a_2 a_3 N(A_u=a_1, A_v=a_2, A_w=a_3)}{\sum'_{a_1, a_2, a_3} N(A_u=a_1, A_v=a_2, A_w=a_3)},\end{aligned}$$

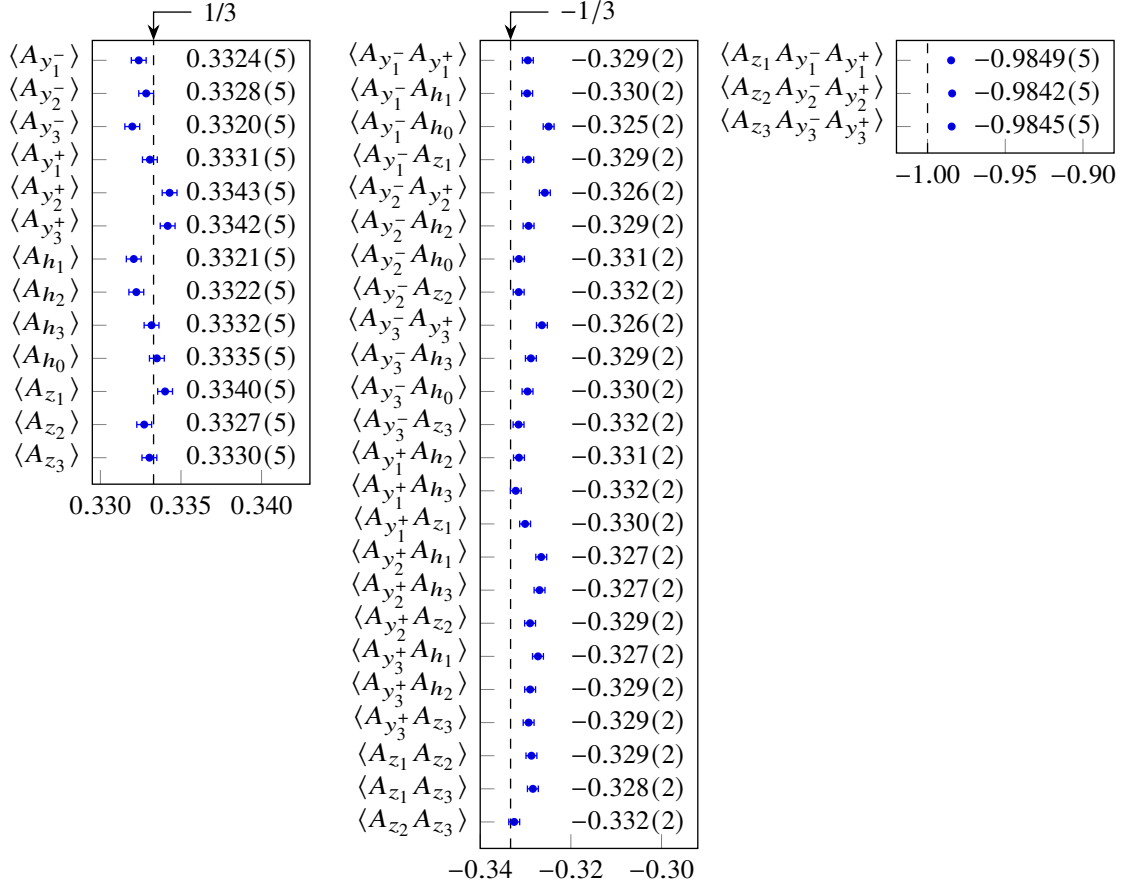
where  $\sum'$  additionally sums over all permutations of the argument list of  $N$ , i.e. the measurement order. Fig. 7.6 displays the measurement results, together with the theoretical predictions for a random input state.

Substituting the obtained values into the Yu-Oh witnesses in Eq. (6.59) and Eq. (6.65) we find

$$\langle \chi_{\text{YO}} \rangle = 8.279(4) \quad \text{and} \quad \langle \chi_{\text{opt}} \rangle = 27.357(11).$$

Our results violate the original Yu-Oh inequality (Eq. 6.4.1) by 69 standard deviations, and the optimal Yu-Oh inequality (Eq. 6.4.2) by 214 standard deviations. Note however that these deviations are solely based on statistical uncertainties, which are small due to a large number of measurements in the complete dataset. While taking more data can decrease the uncertainty further, it does not increase the confidence that a deterministic NC model of reality has been disproved, since that question hinges predominantly on the issue of measurement compatibility and sharpness, as argued in Sec. 6.5. Consequently, we turn our attention to verifying those assumptions.





**Figure 7.6:** Experimental results for expectation values that enter the Yu-Oh witnesses in Eq. (6.59) and Eq. (6.65). Error bars reflect shot noise; dashed lines represent values predicted by quantum mechanics for a random input state.


### 7.3.2. Assumption checks

The Yu-Oh test is based on the assumption that the correlations are measured for perfectly compatible observables, and that individual measurements are perfectly sharp. Both of those assumptions are violated by the data.

Sharp measurements return the same outcome when performed multiple times in a row. To quantify how well this is realized experimentally, we calculate the probability

$$R_u = \frac{\sum_{a_1} N(A_u=a_1, A_u=a_1)}{N(A_u, A_u)}$$

for obtaining the same outcomes in two consecutive measurements of observable  $A_u$ . Ideally, this probability would be 1 for all observables. Instead, we experimentally find them to be lower with deviations on the order of  $2 \times 10^{-3}$ , as shown in Fig. 7.7.



**Figure 7.7:** Measure of the repeatability of our measurements calculated according to Eq. (7.3.2). The ideal value is 1.

The compatibility violation can be seen in Fig. 7.6, where the correlators  $\langle A_v A_u \rangle$  systematically deviate from the expected value of  $-1/3$ . This is an indicator of imperfect compatibility between  $M_v$  and  $M_u$ , which leads to a small number of erroneous outcomes  $N(A_v = -1, A_u = -1)$ . Averaging over all supposedly compatible measurements pairs  $(u, v)$ , we find an error probability of

$$\frac{N(A_v = -1, A_u = -1)}{N(A_v = -1)} = 37(11) \times 10^{-4},$$

in agreement with typical pi-pulse errors at the Doppler temperature (Eq. 7.4). The quantity in brackets represents the standard deviation of the all the 48 observable pairs.

Despite the imperfect compatibility, we do not observe significant *context signalling* in the data. Suppose that measurement  $M_u$  is supposedly compatible with both  $M_v$  and  $M_w$ . While incompatibility means that the presence of either  $M_v$  and  $M_w$  affects the outcome of  $M_u$  when performed before it, context signalling would mean that the outcome of  $M_u$  is affected by whether  $M_v$  or  $M_w$  is performed before. We find no evidence of context signalling in our dataset, and further details can be found in [Leu17, Supp. Mat.].

The measured deviations from the idealised assumptions should affect the evaluated witness and, in turn, determine the true degree of nonclassicality in the correlations. Alas, while some penalisation techniques are known in the KCBS scenario, we are not aware of any standard method to account for these imperfections when evaluating the Yu-Oh witnesses (or any other state-independent NC inequality). Consequently, we are only able to quantify the degree of error, and we hope future theoretical developments will allow for these corrections to be included in our conclusions.

## 7.4. KCBS experiment

This section discusses the experiments showing the KCBS inequality violation with a single  $^{40}\text{Ca}^+$  ion (see Sec. 7.4.4 for a comparison with previous demonstrations). Summarising Sec. 6.3, the KCBS inequality in Eq. (6.47) can be violated by correlation measurements of an initial state  $|0\rangle$  along five directions given by

$$|\psi_i\rangle = U_i|0\rangle \quad \text{with} \quad U_i = R_z^{2(i-1)}(\frac{2\pi}{5})R_y(\theta).$$

where  $i = 1, \dots, 5$  and

$$\theta = \theta_5 = \arctan\left(\frac{2}{\sqrt{1+\sqrt{5}}}\right) \approx 48^\circ$$

We refer to these states as the pentagon stages due to their geometrical interpretation (Fig. 6.2).

### 7.4.1. Measurement construction

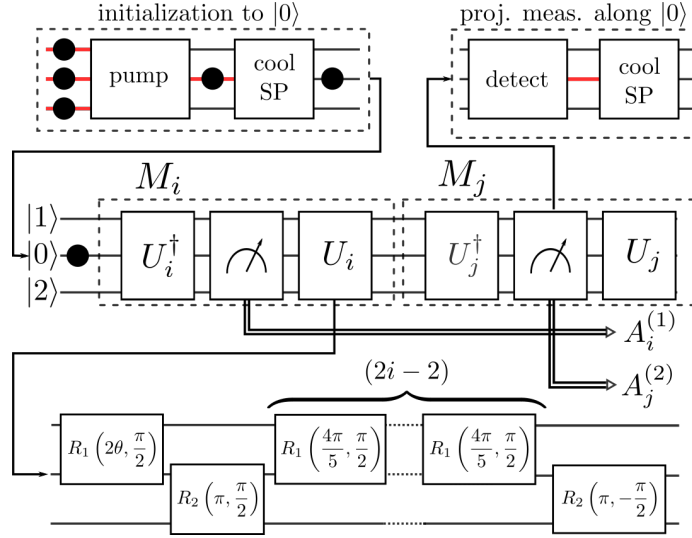
Each unitary  $U_i$  could be decomposed into two rotations as in Sec. 7.1.1. However, in the KCBS experiment we deliberately employed a less efficient pulse decomposition. This was done in order to be able to easily scan the degree of incompatibility between the measurements. Following Eq. (7.2), we implement  $U_i$  as a sequence of  $2i + 1$  coherent rotations:

$$U_i = R_2\left(\pi, \frac{\pi}{2}\right) R_1^{2i-2}\left(\frac{4\pi}{5}, \frac{\pi}{2}\right) R_2\left(\pi, -\frac{\pi}{2}\right) R_1\left(2\theta, \frac{\pi}{2}\right). \quad (7.6)$$

The total measurement sequence for a pair of observables  $\hat{M}_i, \hat{M}_j$  is illustrated in Fig. 7.8.

Aside from improved cooling (Sec. 7.1.3), two key points differentiate the individual KCBS measurements from the Yu-Oh measurements. First, individual observables do not require separate calibrations, as the observable index  $i$  is selected by simply changing the number of  $R_1(4\pi/5, \pi/2)$  pulses, rather than by changing individual pulse areas. Thus, the only recurrent pulse area calibration in the KCBS experiment are the  $R_1(4\pi/5, \pi/2)$  calibration, which is performed by repeating it 10 times on an initial state  $|0\rangle$ , and the  $R_2(\pi, \pm\pi)$

## 7. EXPERIMENTAL STUDIES OF QUANTUM CONTEXTUALITY



**Figure 7.8:** Sequential measurement of observables  $M_i$  and  $M_j$ . Red lines illustrate the parts of the sequence where the ion is hot and/or outside the computational basis. In the *pump* stage, the ion is Doppler-cooled and pumped into the  $S_{1/2}$  manifold, which includes state  $|0\rangle$ . In the *detect* stage, the state is projected either onto the  $\{|1\rangle, |2\rangle\}$  manifold (without affecting its motional state), or onto the  $S_{1/2}$  manifold (heating it back to the Doppler limit). In the *cool SP* stage, the  $S_{1/2}$  states are ground-state cooled and pumped into  $|0\rangle$ , while states  $|1\rangle$  and  $|2\rangle$  remain unaffected. Each unitary  $U_i$  is decomposed into a sequence of  $(2i + 1)$  coherent rotations on  $|0\rangle \leftrightarrow |1\rangle$  and  $|0\rangle \leftrightarrow |2\rangle$  transitions. Every sequence produces outcomes  $A_i^{(1)}$  and  $A_j^{(2)}$ , from which we calculate the correlator  $A_i^{(1)} A_j^{(2)}$ .

calibration. Second, instead of performing the measurement for a single observable set, we run it for a range of values of  $\theta$ . By scanning  $\theta$  around the theoretical compatibility point of  $\theta_5 \approx 48^\circ$ , we can assess the effect of measurement incompatibility on the experiment.

### 7.4.2. Measurement incompatibility

The traditional KCBS inequality is only defined for  $\theta = \theta_5$ , where we find (Eq. (6.47)):

$$S_5(\theta_5) = S_5^\pm(\theta_5) = \sum_{i=1}^5 \langle A_i^{(1)} A_{i\pm 1}^{(2)} \rangle \geq \bar{S}_5^{\text{NC}} = -3. \quad (7.7)$$

The  $\pm$  terms highlight that the witness can be defined in terms of measurements performed either in the “normal order” ( $S_5^+$ ) or the “reverse order” ( $S_5^-$ ), and we used the superscripts to clarify that  $A_i^{(1)}$  is measured first, and  $A_{i\pm 1}^{(2)}$  is measured second. The impossibility of perfect calibration of  $\theta = \theta_5$  creates a compatibility loophole in the experiment (Sec. 6.5). In order to constrain the influence of miscalibrations, we follow an approach introduced in [Kuj15] within the contextuality-by-default framework (Sec. 6.5.2). We quantify the

influence of incompatibility by the degree of signaling  $\epsilon_i$ , which is defined as:

$$\epsilon_i = |\langle A_i^{(1)} \rangle - \langle A_i^{(2)} \rangle| \quad (7.8)$$

where  $\langle A_i^{(1)} \rangle$  is the average outcome of measurement  $M_i$  when it is performed as the first measurement in a sequence, and  $\langle A_i^{(2)} \rangle$  is the outcome when  $M_i$  is the second measurement, i.e. following a measurement of  $M_{i\pm 1}$ . The extended KCBS inequality is then given by

$$S_5^{(\text{ext})}(\theta) = \sum_{i=1}^5 \langle A_i^{(1)} A_{i\pm 1}^{(2)} \rangle + \sum_{i=1}^5 \epsilon_i \geq \bar{S}_5^{\text{NC}} = -3, \quad (7.9)$$

where once again there are two possible measurement orders.

We are now in a position to evaluate the witnesses analytically for any given  $\theta$ . Suppose measurements  $M_i$  and  $M_{i\pm 1}$  are conducted in sequence. Without assuming compatibility, we can use Eqs. (6.15), (6.16) and (6.17) to evaluate the measurement outcomes and the measurement correlator. The calculation gives

$$\begin{aligned} \langle A_i^{(1)} \rangle &= \cos(2\theta) \\ \langle A_{i\pm 1}^{(2)} \rangle &= \cos(2\theta) + \frac{1}{16}(5 - \sqrt{5} + 5(3 + \sqrt{5}) \cos(2\theta)) \sin^2(2\theta) \\ \langle A_i^{(1)} A_{i\pm 1}^{(2)} \rangle &= \frac{1}{8}(3 - \sqrt{5} + (5 + \sqrt{5}) \cos(4\theta)) \end{aligned}$$

Since the results are independent of  $i$ , we can evaluate the signaling term as


$$\epsilon_i = |\langle A_i^{(1)} \rangle - \langle A_i^{(2)} \rangle| = |\langle A_i^{(1)} \rangle - \langle A_{i\pm 1}^{(2)} \rangle|$$

and the witnesses as:

$$\begin{aligned} S_5(\theta) &= 5 \langle A_i^{(1)} A_{i\pm 1}^{(2)} \rangle \\ S_5^{(\text{ext})}(\theta) &= 5 \langle A_i^{(1)} A_{i\pm 1}^{(2)} \rangle + 5\epsilon_i \end{aligned}$$

These expressions are plotted in Fig. 7.9. The graph illustrates a number of important features. The extended witness exhibits nonclassicality for a finite range of opening angles  $\theta \approx (44.5, 50.5)$  deg. For  $\theta = \theta_5$  we indeed find  $\epsilon_i = 0$  and  $S_5^{(\text{ext})} = S_5 \approx -3.944$ , while for all other values of  $\theta$  result in a lower degree of violation.

Interestingly, the minimum of  $S_5(\theta)$  is shifted from  $\theta = \theta_5$  to  $\theta = \pi/4$ , where  $S_5 = -5(1 + \sqrt{5})/4 \approx -4.045$ . This highlights the reason behind the need for incompatibility penalisation since measurement miscalibration can *increase* the amount of (apparent) classicality violation in the traditional witness. For the extended witness, on the other hand, the minimum of  $S_5^{(\text{ext})}(\theta)$  is the compatibility point. We also verified through an explicit calculation that no single pulse angle of phase miscalibration can decrease the measured value of  $S_5^{(\text{ext})}$ , although we have not confirmed if that minimum is global with respect to



**Figure 7.9:** Predicted values of  $S_5$  (Eq. ((7.7))) and  $S_5^{(\text{ext})}$  (Eq. ((7.9))) as a function of  $\theta$ . The hashed region above  $-3$  shows the space where  $S_5^{(\text{ext})}$  is consistent with deterministic NC models. The inset shows how a finite number of experiments (here  $5 \times 10^4$  per data point) leads to a necessary deviation from theory around  $\theta_5$  (dashed line).

all parameters.

### Influence of shot noise

Aside from systematic effects, some amount of signalling is expected purely due to shot-noise-limited statistics. Due to the absolute value function in the definition of  $\epsilon_i$  (Eq. (7.8)), finite sample size creates a systematic difference between the experimentally measured sample mean  $\langle \epsilon_i \rangle$  and the population mean. This results in an increase of  $S_5^{(\text{ext})}$  for any finite number of experiments which is particularly pronounced around  $\epsilon_i = 0$ . Consider an experimental run with  $N$  measurement repetitions. The measured outcomes follow a normal distribution

$$Y_{\langle A_i \rangle} \propto \mathcal{N}(\langle A_i \rangle, \sigma_{A_i}^2),$$

where  $\mathcal{N}(\mu, \sigma)$  denotes a normal distribution with mean  $\mu$  and variance  $\sigma^2$ . Eq. (2.24) implies that the variance is given by

$$\sigma_{A_i}^2 = \frac{1 - \langle A_i \rangle^2}{N}$$

where we assumed binomial statistics, which apply when  $\langle A_i \rangle$  is not very close to  $\pm 1$ . The difference between the measurement averages is distributed as:

$$Y_{\langle A_i \rangle - \langle A_{i+1} \rangle} \propto \mathcal{N}(\langle A_i \rangle - \langle A_{i+1} \rangle, \sigma_{A_i}^2 + \sigma_{A_{i+1}}^2)$$

and thus the signaling term is distributed according to:

$$Y_{\epsilon_i} = \mathcal{F}(\langle A_i \rangle - \langle A_{i+1} \rangle, \sigma_{A_i}^2 + \sigma_{A_{i+1}}^2)$$

where  $\mathcal{F}(\mu, \sigma^2)$  is the so-called “folded normal distribution”, obtained by taking the absolute value of a Gaussian with mean  $\mu$  and variance  $\sigma^2$ . The mean of a folded normal distribution is [Leo61]:

$$\mu_F(\mu, \sigma) = \sigma \sqrt{\frac{2}{\pi}} \exp\left(\frac{-\mu^2}{2\sigma^2}\right) - \mu \operatorname{erf}\left(\frac{-\mu}{\sqrt{2}\sigma}\right),$$

where  $\operatorname{erf}(x)$  is the error function. This mean represents an expectation value  $\langle \epsilon_i \rangle$  of  $\epsilon_i$  in an experiment with  $N$  repetitions:

$$\langle \epsilon_i \rangle = \mu_F(\epsilon_i, \sigma_{A_i}^2 + \sigma_{A_{i+1}}^2)$$

and the expectation value of the extended witness is given by

$$\langle S_5^{(\text{ext})}(\theta) \rangle = 5\langle A_i^{(1)} A_{i\pm 1}^{(2)} \rangle + 5\langle \epsilon_i \rangle.$$

This is shown as a red dashed line in Fig. 7.9 for  $N = 10,000$ . Indeed, the effect is usually negligible, except for  $\theta \approx \theta_5$ , where we expect a gap of

$$\langle S_5^{(\text{ext})}(\theta) \rangle - S_5^{(\text{ext})}(\theta) \approx 10 \sin(\theta) \sqrt{\frac{2}{\pi N}},$$

which evaluates to

$$\langle S_5^{(\text{ext})}(\theta_5) \rangle - S_5^{(\text{ext})}(\theta_5) \approx 5.93 \sqrt{\frac{1}{N}}.$$

For  $N = 10,000$  experiments per observable pair the gap is 0.06, leading to the minimum observable witness value of

$$\langle S_5^{(\text{ext})}(\theta_5) \rangle \approx -3.885.$$

### 7.4.3. KCBS inequality violation

For each opening angle setting  $\theta_{\text{set}}$  we perform a correlation measurement along a pair of directions ( $|\psi_i\rangle, |\psi_j\rangle$ ), with results ( $A_i^{(1)} = \pm 1, A_j^{(2)} = \pm 1$ ) and correlation  $A_i^{(1)} A_j^{(2)}$ . Each experiment is repeated 10,000 times, and the average results ( $\langle A_i^{(1)} \rangle, \langle A_j^{(2)} \rangle, \langle A_i^{(1)} A_j^{(2)} \rangle$ ) are extracted. We collect data for 5 observable pairs in the normal order, and 5 in the reverse order, and evaluate the witnesses using Eq. (7.7) and Eq. (7.9). We also estimate the actual

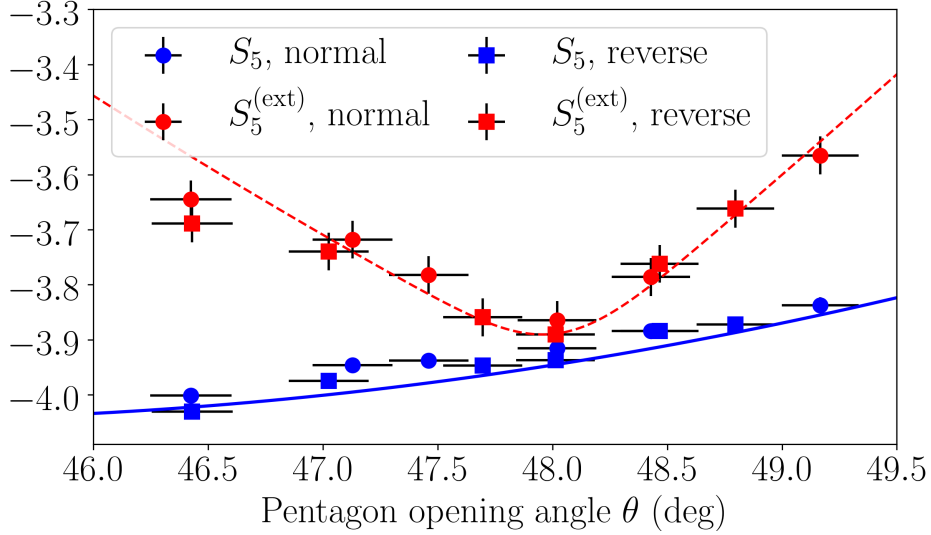
## 7. EXPERIMENTAL STUDIES OF QUANTUM CONTEXTUALITY

pentagon opening angle  $\theta$  using

$$\theta = \frac{1}{2} \arccos \left( \sum_{i=1}^5 \langle A_i^{(1)} \rangle / 5 \right).$$

For each average ( $\langle A_i^{(1)} \rangle, \langle A_j^{(2)} \rangle, \langle A_i^{(1)} A_j^{(2)} \rangle$ ) we extract the sample standard deviation from the data and use it to compute the standard error in the mean. We then propagate the uncertainty to  $S_5, \theta, \epsilon$  and  $S_5^{(\text{ext})}$  assuming independent errors. We note that  $\epsilon$  is defined as necessarily positive and its distribution is non-Gaussian, hence the standard error of  $\epsilon$  cannot always be treated as a confidence interval.

The measured witnesses  $S_5(\theta)$  and  $S_5^{(\text{ext})}(\theta)$  are displayed in Fig. 7.10, together with theoretical expectations for ideal shot noise-limited measurements. The results of  $S_5(\theta)$  exhibit



**Figure 7.10:** Results of the KCBS measurement as a function of the opening angle  $\theta$ . Each data point results from 10,000 measurements on each of the five correlators  $\langle A_i A_j \rangle$ , either in normal ( $j = i + 1$ ) or reverse order ( $j = i - 1$ ). Blue and dashed red lines represent theoretical expectations for  $S_5$  and  $\langle S_5^{(\text{ext})} \rangle$  respectively. Note that all measurements of  $S_5^{(\text{ext})}$  violate the NC bound of  $S_5^{\text{NC}} = -3$ . Error bars here show the standard error in the mean, with sample standard deviation obtained directly from the data.

a systematic shift of 1.6 standard deviations from the ideal QM prediction, attributable mainly to qutrit rotation errors (Sec. 7.1.1). For the extended witness  $S_5^{(\text{ext})}(\theta)$ , statistical errors dominate. Table 7.2 lists the full the results for the value of  $\theta$  measured to be closest to  $\theta_5 \approx 48^\circ$  in each scan. The data point closest to compatibility gives:

$S_5 = -3.915(14)$	$S_5^{(\text{ext})} = -3.864(34)$	normal order
$S_5 = -3.937(14)$	$S_5^{(\text{ext})} = -3.890(34)$	reverse order



**Table 7.2:** Experimental results for the KCBS experiment for the points closest to the compatibility angle  $\theta = \theta_5$  in both normal and reverse order in Fig. 7.10.

Order	$i$	$j$	$\langle A_i \rangle$	$\langle A_j \rangle$	$\langle A_i A_j \rangle$
Ideal			$\approx -0.105$	$\approx -0.105$	$\approx -0.788$
Ideal total			$S_5 \approx -3.944, S_5^{(\text{ext})} \approx -3.888$		
Normal	1	2	-0.106(10)	-0.107(10)	-0.786(6)
	2	3	-0.111(10)	-0.092(10)	-0.793(6)
	3	4	-0.107(10)	-0.112(10)	-0.775(6)
	4	5	-0.102(10)	-0.107(10)	-0.787(6)
	5	1	-0.100(10)	-0.121(10)	-0.774(6)
	Total		$S_5 = -3.915(14), S_5^{(\text{ext})} = -3.864(34)$		
Reverse	1	2	-0.113(10)	-0.096(10)	-0.786(6)
	2	3	-0.111(10)	-0.101(10)	-0.787(6)
	3	4	-0.106(10)	-0.103(10)	-0.784(6)
	4	5	-0.093(10)	-0.118(10)	-0.783(6)
	5	1	-0.102(10)	-0.097(10)	-0.798(6)
	Total		$S_5 = -3.937(14), S_5^{(\text{ext})} = -3.890(34)$		

which violates the KCBS Ineq. (7.7) by 65 (67) standard deviations in the normal (reverse) order and the extended KCBS Ineq. (7.9) by 25 (26) standard deviations.

#### 7.4.4. Comparison with other experiments

The experimentally obtained values of  $\langle S_5(\theta_5) \rangle$  and  $\langle S_5^{(\text{ext})}(\theta_5) \rangle$  are the lowest of any KCBS test to date. In addition, we believe we improve on the state-of-the-art by simultaneously addressing a range of issues left open by other experiments, as listed in Tab. 7.3.

Aside from being among the closest to QM predictions, our analysis is the only one that we are aware of that systematically characterizes signalling and is in agreement with theoretical expectations. We have calculated the signalling the different experiments would have incurred wherever the data was made available by the authors (second-to-last column). The meanings of the different comments in the last column are:

- *Detection loophole.* Experiments with photons suffer from significant detector inefficiencies and setup losses. Consequently, the results assume that the registered events form an unbiased sample of all the events. The loophole does not apply when measurements are conducted with high fidelity.
- *Simultaneous measurements.* When pairs of measurements are conducted simultaneously, it is difficult to establish an operational definition of an individual measurement [Jer16]. It has been argued that this so-called “individual existence loophole” can be addressed by performing measurements in a sequence [Cab09].
- *Six observables.* Certain experimental arrangements may not allow for performing

## 7. EXPERIMENTAL STUDIES OF QUANTUM CONTEXTUALITY

**Table 7.3:** Experimental results of previous KCBS tests. Comments in the last column are discussed in the text. The results marked by an asterisk represent our own analysis of the results table from [Um13] and private communications with the authors.

Reference	Platform	Saturation of QM limit $(S_5 - \bar{S}_5^{\text{NC}})/$ $(\bar{S}_5^{\text{QM}} - \bar{S}_5^{\text{NC}})$	Signaling $\sum_{i=1}^5 \epsilon_i /$ $(\bar{S}_5^{\text{QM}} - \bar{S}_5^{\text{NC}})$	Comments
[Lap11]	Photons	0.947(6)	0.08(3)	Detection loophole Simultaneous measure- ments Six observables
[Ahr13]	Photons	0.53(11) (normal order) 0.95(11) (reverse order)	No data	Detection loophole Order dependence
[Den13]	Photons	0.977(11) (normal order) 0.956(26) (reverse order)	0.267  0.291	Detection loophole Simultaneous measure- ments Large signaling
[Um13]	Yb ion	0.589(24)*	0.119(24)*	Six observables Non-projective measure- ments
[Jer16]	Supercond. circuits	0.520(1) (normal order) 0.541(1) (reverse order)	0.379(2)  0.379(2)	Large signaling
<b>This work</b>	Ca ion	<b>0.969(14)</b> (normal order) <b>0.992(14)</b> (reverse order)	<b>0.054(31)</b>  <b>0.050(31)</b>	

the same measurement in the same way in every context. In order to circumvent this problem, the KCBS inequality is extended to a six-observable inequality. However, the validity of this approach has been challenged by some authors [Ahr13], although the issue was contested [Lap13].

- *Order dependence.* For projective measurement, the order of measurement should not influence the correlator (Sec. 6.1.3). Significant dependence on the order indicates large measurement noise or uncontrolled drifts, questioning the sharpness assumption.
- *Large signaling.* When signalling is large compared to the shot-noise limit, the experiment is far from the ideal assumption of compatible sharp measurements.
- *Non-projective measurements.* In Ref. [Um13] measurements were conducted in a sequence but were not projective by design. Specifically, the post-measurement state, conditioned on photon detection, is a mixture of qutrit basis states. This precludes the implementation of sharp measurements.

The largest outstanding issue in our experiment (which is also present in all the other experiments) is the measurement sharpness assumption. While measurement repeatability errors can be quantified, they are not penalised in the extended witness. Recent theoretical work overcomes this issue in the operational framework [Kun19], but it remains to be clarified in the other frameworks.

## 7.5. Generalised KCBS experiment

The generalised KCBS inequality in Eq. 6.56 can be violated by correlation measurements of an initial state  $|0\rangle$  along  $N$  directions given by

$$|\psi_i\rangle = U_i|0\rangle \quad \text{with} \quad U_i = R_z^{(i-1)(N-1)/2}(\frac{2\pi}{N})R_y(\theta).$$

where  $N \geq 5$  is odd and

$$\theta = \theta_N = \arccos \sqrt{\frac{\cos(\frac{\pi}{N})}{1 + \cos(\frac{\pi}{N})}}.$$

For an ideal experiment, the contextual content (Eq. (6.58))  $q \rightarrow 1$  as  $N \rightarrow \infty$ , indicating fully contextual correlations.

### 7.5.1. Ray construction

Similarly to Eq. (7.6), we construct  $U_i$  as a sequence of  $3 + (i-1)(N-1)/2$  coherent rotations:

$$U_i = R_2\left(\pi, \frac{\pi}{2}\right) R_1^{(i-1)(N-1)/2}\left(\frac{4\pi}{5}, \frac{\pi}{2}\right) R_2\left(\pi, -\frac{\pi}{2}\right) R_1\left(2\theta, \frac{\pi}{2}\right).$$

## 7. EXPERIMENTAL STUDIES OF QUANTUM CONTEXTUALITY

However, the inefficiency of this construction becomes particularly noticeable for large  $N$ , where the typical number of rotations per sequence grows  $\propto N^2$ . In order to decrease the runtime while retaining the calibration simplicity, we concatenate the pulse sequence performed between the two sequential measurements. Recall that for the measurement of an observable sequence  $\hat{M}_i, \hat{M}_{i+1}$ , the unitary implemented between the detection steps is  $U_{i+1}U_i$  (Sec. 7.1.3), which can be evaluated as

$$U_{i+1}^\dagger U_i = R_y(-\theta) R_z^{(i+1)(N-1)/2} \left(-\frac{2\pi}{N}\right) R_z^{(i-1)(N-1)/2} \left(\frac{2\pi}{N}\right) R_y(\theta)$$

which simplifies to

$$U_{i+1}^\dagger U_i = R_y(-\theta) R_z^{-(N-1)} \left(\frac{2\pi}{N}\right) R_y(\theta)$$

allowing for a pulse decomposition that only involves  $N + 3$  pulses

$$U_{i+1}^\dagger U_i = R_1 \left(2\theta, \frac{3\pi}{2}\right) R_2 \left(\pi, \frac{\pi}{2}\right) R_1^{(N-1)} \left(\frac{4\pi}{N}, \frac{3\pi}{2}\right) R_2 \left(\pi, -\frac{\pi}{2}\right) R_1 \left(2\theta, \frac{\pi}{2}\right).$$

and  $R_1 \left(\frac{4\pi}{N}, \frac{3\pi}{2}\right)$  is calibrated by applying it  $N$  times to an initial state  $|0\rangle$ . This construction leads to relevant time and infidelity reductions for large  $N$ , but it opens the “individual existence loophole”, discussed in Sec. 7.4.4 as a weakness of experiments with simultaneous measurements.

### 7.5.2. Witness construction

In analogy to Sec. 7.4.2, we define the traditional and the extended inequality as

$$S_N = \sum_{i=1}^N \langle A_i A_{i+1} \rangle \geq \bar{S}_N^{\text{NC}} = -N + 2$$

$$S_N^{(\text{ext})} = \sum_{i=1}^N \langle A_i^{(1)} A_{i+1}^{(2)} \rangle + \sum_{i=1}^N \epsilon_i \geq \bar{S}_N^{\text{NC}} = -N + 2,$$


where  $\epsilon_i = |\langle A_i^{(1)} \rangle - \langle A_i^{(2)} \rangle|$  and  $S_N^{(\text{ext})}$  reduces to  $S_N$  when  $\theta = \theta_N$ . Using Eq. (6.58), we find the contextual content can be calculated using

$$q_N = \frac{S_N + N - 2}{2}.$$

although I do not know of any work which extends the definition  $q$  to measurements with non-zero signaling.

### 7.5.3. Experimental results

We performed the experiments with  $N$  increasing from  $N = 5$  to  $N = 121$ . Each experiment is conducted in analogy to the KCBS test, with the pentagon opening angle calibrated to  $\theta \approx \theta_N$ . All measurements were performed in the normal order, with each of the  $N$  compatible observable pairs measured 10,000 times.



**Figure 7.11:** Results of measurements of  $N$ -cycle witnesses. Solid and dashed lines show QM expectations for relevant quantities (red dashed line includes shot noise), and hashed regions below 0 correspond to results explainable by classical models. The top plot illustrates the fractional gap between quantum and classical witnesses, which decreases for large  $N$ . Our data shows contextuality up to  $N = 101$  ( $N = 61$ ) for  $S_N$  ( $S_N^{(\text{ext})}$ ). The bottom plot shows the calculated contextual fraction. Ideally the system becomes fully contextual at  $N \rightarrow \infty$ , but experimental imperfections lead to  $\text{CF} < 0$  for large enough  $N$ . We measure  $\text{CF}_{31} = 0.800(4)$ . Error bars are generally smaller than the point size.

The measurement results are shown in Fig. 7.11. We identify stronger-than-classical correlations in all prepared scenarios up to  $N = 101$  for the traditional witness, and up to  $N = 61$  for the extended witness. Beyond that, our results are consistent with NC models. The largest measured contextual content is  $q_{61} = 0.800(4)$ . The exact results for each observable are listed in [Mal17, Supp.Mat. I].

To our best knowledge, these results show contextuality in a system with the largest

number of observables (101) of any experiment reported up to this date. This study complements previous tests of chained Bell inequalities, which have measured contextual contents of as large as  $q_{36} = 0.874(1)$  with entangled photons [Chr15] and  $q_{18} = 0.704(12)$  with entangled trapped ions [Tan17].

### Quantum-to-classical transition?

While QM predicts that  $q_N$  will approach unity as  $N$  increases, any finite error rate per measurement will cause  $q_N < 1$  for sufficiently large  $N$ . One way of looking at this (admittedly a counter-intuitive one!) is as a quantum-to-classical transition [Kof07; Jeo09]: whereas the correlations for  $N \leq 101$  cannot have deterministic NC models, the correlations we measure for 121 could have been produced classically. Assuming a constant error rate per measurement, this transition towards classicality does not originate from increased interactions with an environment (measured coherence times significantly exceed the sequence lengths for any  $N$ ) and is not due to an increase in the macroscopicity of the system (the Hilbert-space dimension is always 3). Instead, classical expectations emerge when the measurement directions are spaced closely enough that they cannot be well-resolved from each other. This possibility was first discussed in [Kof07], where it is proven that measurement imperfections in an otherwise perfectly quantum system can give rise to classical behaviour. It was later shown that measurement errors do not lead to a quantum-to-classical transition if the final projection is coarsened, whereas they do when measurement references are coarsened [Jeo14], i.e. when unitary operations are non-ideal, which is consistent with our observations.

## 8. Entanglement with integrated waveguides

This chapter discusses two-ion quantum logic with integrated waveguides. It is mainly focused on the basic two-ion operations and the calibration and performance of the MS gate in the trap #4. This chapter expands on the results published in [Meh20]. Sec. 8.1 discusses the basic tools and benchmarks of two-ion trapping, control and readout. Sec. 8.2 discusses different types of substrate charging observed in the integrated-waveguide traps, and how the gate is calibrated. Sec. 8.3 presents the Bell-state fidelity measurement and an analysis of possible error sources. All the results presented in this chapter are in agreement with those reported in [Meh20], with small discrepancies in the exact parameter values due to slight differences in the analysis methods, rounding errors and dataset selection.

While the highest-fidelity Bell states were generated in the trap #4, two-qubit lifetime deterioration that occurred not long after those measurements prohibited us from extensively exploring certain noise sources. Therefore, to complement the published results, this chapter presents selected additional data from trap #5. These are used to complete the discussions of axial stretch mode heating rates, light-induced grating coupler charging, gate calibration methods and gate error sources.


### 8.1. Basic two-ion operations

#### 8.1.1. Trapping two ions

By far the most challenging aspects of performing two-ion experiments in our setup are slow loading and limited lifetimes of the second ion. While single-ion loading is reliable (Sec. 5.1.1), the second ion can almost never be loaded in the single-ion working-well.

The most reproducible way of loading the second ion involves using a dedicated lower-frequency loading-well. Following single-ion micromotion compensation in the working-well (Sec. 5.1.2), we lower the axial frequency to  $\omega \lesssim 2\pi \times 0.5$  MHz, keeping the compensation fields unchanged. Initially, we scan the axial frequency while attempting to trap the second ion, looking for a potential well with a maximal trapping rate. If none is suitable, we adjust the compensation fields hoping to trap. Once an appropriate well is found, we can use it in the future, and typically load the second ion within 5 minutes. However, the trapping well is selected somewhat blindly, and there are days where the second ion doesn't get trapped at all.

Evidently, the main issue is that the trapping probability depends not only on the voltage set, but also on whether there is already an ion trapped in that same well. In an attempt to circumvent this problem, we implemented a method of loading based on ion transport<sup>1</sup> [Vit13]. The first ion is trapped and cooled in the working-well. Then, the ion is displaced by  $\approx 300\mu\text{m}$  along the trap axis, and an empty working-well is created at the position where the first ion was loaded. We wait until a second ion is trapped in the working-well. Following fluorescence detection, the two wells with one ion each are merged into one well with two ions. If the second ion is not trapped within  $\tau = 10\text{sec}$ , we transport the first ion back to the lasers' position for Doppler cooling and repeat the sequence. The value of  $\tau$  is low enough that we do not record significant ion loss probability over hundreds of repetitions. Fig. 8.1 shows a typical fluorescence trace of the transport loading sequence.



**Figure 8.1:** Measured fluorescence counts during a transport loading sequence. The first ion is moved away 7 times over 45 seconds. During that time, a second ion is trapped twice, but only the second attempt to merge the two ions into a single well is successful. A successful crystallisation of two ions in a single well is inferred from the approximate doubling of the fluorescence signal.

The second ion is loaded during the second merge attempt. This was common, as the merge sequence is only successful  $\approx 50\%$  of the time.

For a period of several months, transport loading allowed for reliable and fast creation of two-ion crystals. Alas, all good things come to an end. At a certain point, we began observing significant loss probability during ion transport, which soon became near-deterministic. We noticed that it was possible to modify the compensation fields during transport to reduce the loss, but the calibrations were tedious and short-lived. Since the lasers were only aligned at the working-well, the compensation errors during transport could only be inferred by measuring the ion returning to the working-well following transport [Kau20]. However, the waveform control was not synchronous with the experiment, but was done through a slow serial communication interface (Sec. 3.4.4). As a result, we were not able to perform

<sup>1</sup> Thanks to C. Ballance for this idea




systematic measurements of transport-induced motional excitation. Instead, the signal was binary: either the ion was lost at the end of the sequence, or not. As expected, that was not a very pleasant way to take data.

As a result, we returned to the first loading method for the time being. In the meantime, we prepared to upgrade the trap voltage sources and their interface to allow for fast and synchronous control. The new sources were installed by C. Mordini during the writing of this thesis, and we plan to use the transport loading method once again in the near future.

Following the trapping of the second ion, the lifetime of the two-ion crystal is usually limited to approximately 1-2 hours, with frequent losses within the first 10 minutes. This is comparable to two-ion lifetimes in room-temperature SETs and is significantly worse than typical values obtained in other cryogenic SET experiments [All12; Ant09]. The main suspect is that the vacuum pressure is anomalously high, perhaps due to the oven, the vacuum seals (some KF flanges instead of all CF) or virtual leaks. In the future, it would be interesting to more directly estimate the vacuum pressure through double-well measurements with a single ion [Aik20; Cla21].

### 8.1.2. Motional modes and cooling

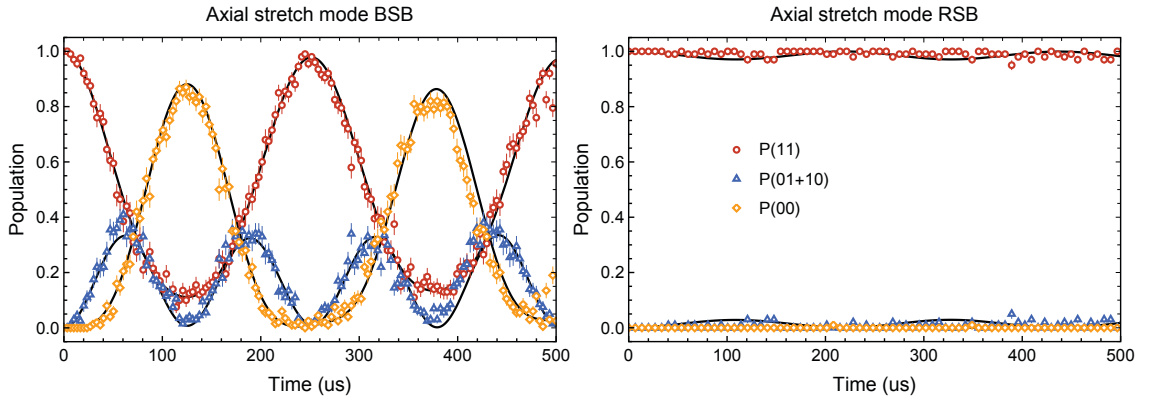


**Figure 8.2:** Blue-sideband spectrum of a two-ion chain. The ratio of the axial stretch (STR) mode frequency and the axial centre-of-mass (COM) mode frequency is  $\sqrt{3}$ .

A two-ion  $^{40}\text{Ca}^+$ - $^{40}\text{Ca}^+$  chain has six normal modes of motion. In addition to the three centre-of-mass (COM) oscillation modes of a single ion, we now also find three modes where the centre-of-mass remains stationary, but the ions oscillate out of phase with each other. The axial out-of-phase mode is referred to as the stretch mode, and for a harmonic trap its frequency is shifted up from the axial COM frequency by a factor of  $\sqrt{3}$ . The out-of-phase radial modes are referred to as the rocking modes. The resulting two-ion sideband spectrum is shown in Fig. 8.2.

The MS gates described in Sec. 8.3 are mediated via the axial stretch mode. The prerequisite cooling sequence begins by Doppler cooling all the motional modes and EIT cooling the axial modes. We typically find that the EIT cooling settings calibrated with a single ion (Sec. 2.5.3) are appropriate to cool both of the axial modes to  $\bar{n} = 0.4 - 0.5$ , comparable to the single-ion limit. We then execute a pulsed sideband cooling sequence on the axial stretch mode. Each sideband cooling loop involves a BSB pulse on the optical qubit transition followed by an optical pumping pulse on a dipole transition (Sec. 2.5.4). The BSB pulse time is ramped from  $t = 80 \mu\text{s}$  to  $t = 100 \mu\text{s}$ , and the whole sequence takes between 3 to 4 loops to converge.

We can estimate the steady-state gate mode temperature with the sideband fitting method (Sec. 5.6.1). Fig. 8.3 shows the time-traces of the ions' evolution under the drive resonant with the axial stretch mode BSB (left) and RSB (right). The solid curves are the analytical



**Figure 8.3:** Sideband oscillations on the axial stretch mode at  $\omega = 2\pi \times 2.14 \text{ MHz}$  after 3 cycles of sideband cooling. The solid traces correspond to indicate final mode occupation of  $\bar{n} = 0.03$ . The expression was fitted by hand to the BSB trace, and the RSB trace is derived from that fit.

expressions for a symmetric two-ion sideband drive derived in [Kin98; Hom06], implying a thermal occupation of  $\bar{n} = 0.03$ . The temperature could likely be lowered further by smoothly ramping the sideband pulses on- and off (Sec. 2.3.3)


### 8.1.3. Optical qubit Rabi oscillations

Applying a  $\lambda = 729 \text{ nm}$  laser resonant with the optical qubit transition  $|S_{1/2}, m_J = -1/2\rangle \leftrightarrow |D_{5/2}, m_J = -1/2\rangle$  drives Rabi oscillations of both ions simultaneously. However, unless the ions are positioned precisely in the centre of the beam, the contributions of each ion will oscillate at a different rates, producing a beat pattern. Consider two ions prepared initially in  $|11\rangle$  evolving under a carrier drive with Rabi frequencies  $\Omega_1$  and  $\Omega_2$ . Eq. (2.20)

implies that the two-ion populations at time  $t$  are given by:

$$\begin{aligned} P(11) &= \cos(\Omega_1 t/2)^2 \cos(\Omega_2 t/2)^2 \\ P(01 + 10) &= \sin(\Omega_1 t/2)^2 \cos(\Omega_2 t/2)^2 + \sin(\Omega_2 t/2)^2 \cos(\Omega_1 t/2)^2 \\ P(00) &= \sin(\Omega_1 t/2)^2 \sin(\Omega_2 t/2)^2 \end{aligned} \quad (8.1)$$


The position of equal illumination ( $\Omega_1 = \Omega_2$ ) can be found by monitoring  $P(11)$  and  $P(01 + 10)$  while translating the ions' centre-of-mass along the trap axis  $x$  by changing the electrode voltages. The desired position corresponds to the locally fastest oscillations, and the lowest peak value of  $P(01 + 10)$ . Fig. 8.4 shows an example calibration, and the optimal position  $x_0 \approx 1.3 \mu\text{m}$  is the symmetry axis of the scan.



**Figure 8.4:** Optical qubit Rabi oscillations of two ions as a function of the ions' position  $x$  along the trap axis, calculated assuming an centre-of-mass frequency of  $2\pi \times 1 \text{ MHz}$ . The ions are prepared in  $|11\rangle$  at time  $t = 0$ . (Left) Probability  $P(11)$  to find the ions in  $|11\rangle$ . (Right) Probability  $P(01 + 10)$  to find one ion in  $|1\rangle$  and one ion in  $|0\rangle$ .

At the position of nominally equal illumination, we can drive a carrier Rabi oscillation to extract the residual Rabi frequency imbalance and noise. The experiment begins with motional mode cooling as described in Sec. 8.1.2. The qubit state is optically pumped to  $|11\rangle$ , and the resulting Rabi oscillations are shown in Fig. 8.5. The small amount of residual Rabi frequency imbalance can be seen as a slow increase in the moving average of  $P(01 + 10)$ . In Fig. 8.5 (right), the mid-point of the oscillations rises from  $P(01 + 10) = 0.25$  to  $P(01 + 10) = 0.29$ , which according to the simulations implies a Rabi-frequency imbalance of

$$|\epsilon| = |1 - \Omega_2/\Omega_1| = 0.01. \quad (8.2)$$



**Figure 8.5:** Two-ion Rabi oscillation at the calibrated position of equal illumination. The ions are cooled as described in Sec. 8.1.2 and prepared in  $|11\rangle$  at time  $t = 0$ . The solid lines show theoretical traces with Rabi frequency imbalance  $\epsilon = 0.01$  and shot-to-shot fluctuation of  $\sigma_\Omega/\langle\Omega\rangle = 1.0 \times 10^{-2}$ . (Left) Probability  $P(11)$  to find the ions in  $|11\rangle$ . (Right) Probability  $P(01 + 10)$  to find one ion in  $|1\rangle$  and one ion in  $|0\rangle$ .

The observed decay of both curves can be accounted for by Rabi frequency fluctuations. Following the finding in Sec. 5.5.2, we assume the Rabi frequency fluctuations to be dominated by the thermal mode occupancy, and hence  $\Omega_1$  and  $\Omega_2$  can be considered constant over any single experimental shot. Furthermore, the noise is assumed to be common mode to both ions, as will be the case for mode temperature errors and polarization fluctuations, but not for ion position fluctuations). This is modelled by setting  $\Omega_2 = \Omega_1(1+\epsilon)$  in Eq. (8.1) and sampling  $\Omega_1$  from a Gaussian distribution with mean  $\langle\Omega\rangle$  and standard deviation  $\sigma_\Omega$ . The manual fit in Fig. 8.5 corresponds to:


$$\sigma_\Omega/\langle\Omega\rangle = 1.0 \times 10^{-2}, \quad (8.3)$$

in agreement with the single-ion results in Sec. 5.5.2.

#### 8.1.4. Two-ion readout

As discussed in Sec. 2.3.6, the two-ion optical qubit readout proceeds analogously to a single-ion readout, but with two thresholds ( $t_1, t_2 > t_1$ ) to separate the three distinguishable states with zero, one and two ions bright respectively. Fig. 8.6 shows the typical shot record for a two-ion experiment, together with fitted Poissonian peaks (top) and estimated detection errors (bottom). The shots histogram is fitted using three Poisson distributions  $DD(n)$ ,  $SD(n)$  and  $SS(n)$  corresponding to zero, one and two bright ions respectively. The fits indicate mean counts of (3.2, 63.5, 123.3) respectively in a detection window of  $t_{\text{det}} = 250 \mu\text{s}$ .

To compute the estimated readout errors, we correct the distributions for the possibility of ion decay during detection. This is done as in Sec. 2.3.6, but with  $SD(n)$  corrected for



**Figure 8.6:** Two-ion readout histograms. (Top) Histogram  $H_n$  of experimental count distribution (black) for a parity scan a Bell state (Sec. 8.3). The fits show  $H(n) = p_{DD}DD(n) + p_{SD}SD(n) + p_{SS}SS(n)$ , where  $DD(n)$ ,  $SD(n)$  and  $SS(n)$  are Poissonian distributions, and  $p_{DD}$ ,  $p_{SD}$  and  $p_{SS}$  are the fitted probabilities. (Bottom) Simulated detection errors for the fitted distributions, corrected for finite optical qubit lifetime. The dashed lines indicate threshold values where the sum of errors is minimised: here,  $t_1 = 32$  and  $t_2 = 90$ .

a finite dark-state lifetime of  $\tau = 1.1$  s and  $DD(n)$  corrected for an effective two-ion dark state lifetime of  $\tau/2 = 0.55$  s. Extending Eq. (5.1), we define four error parameters for two-ion detection depending on the threshold values  $t_1$  and  $t_2$ :

$$\begin{aligned} \epsilon_0 &= 1 - \text{CDF}(DD(n_{t_1})) & \epsilon_1 &= \text{CDF}(SD(n_{t_1})) \\ \epsilon_2 &= 1 - \text{CDF}(SD(n_{t_2})) & \epsilon_3 &= \text{CDF}(SS(n_{t_2})) \end{aligned} \quad (8.4)$$

As shown in Fig. 8.6 (bottom), the detection errors are now dominated by the overlap of the  $SD(n)$  and  $SS(n)$  histograms. To minimise those, we increase the power of the  $\lambda = 397$  nm laser to operate closer to saturation. This increases the background, but allows  $\epsilon_2$  and  $\epsilon_3$  to be decreased to  $\lesssim 10^{-3}$ . The  $|00\rangle$  readout error is  $\epsilon_0 \approx 3 \times 10^{-4}$ , which is approximately twice that for a single-ion since  $|00\rangle$  has half the lifetime of  $|0\rangle$ .

Once measured, the detection errors can be corrected for in post-processing. The relation between the measured probabilities  $P(DD)$ ,  $P(BD + DB)$ ,  $P(BB)$  of recording zero, one and two ions bright respectively, and the true probabilities  $P(00)$ ,  $P(01 + 10)$ ,  $P(11)$  for zero,

one of two ions to be in  $|1\rangle$  respectively, can be described using the following transformation (c.f. Eq. (5.2)):

$$\begin{pmatrix} P(DD) \\ P(BD + DB) \\ P(BB) \end{pmatrix} = \begin{pmatrix} 1 - \epsilon_0 & \epsilon_1 & 0 \\ \epsilon_0 & 1 - \epsilon_1 - \epsilon_2 & \epsilon_3 \\ 0 & \epsilon_2 & 1 - \epsilon_3 \end{pmatrix} \begin{pmatrix} P(00) \\ P(01 + 10) \\ P(11) \end{pmatrix},$$

where we assumed there is no overlap between  $DD(n)$  and  $SS(n)$ . Inverting the matrix and expanding the resulting up to first order in  $\epsilon_i$ ,  $i = 0, 1, 2, 3$  gives:

$$\begin{pmatrix} P(00) \\ P(01 + 10) \\ P(11) \end{pmatrix} = \begin{pmatrix} 1 + \epsilon_0 & -\epsilon_1 & 0 \\ -\epsilon_0 & 1 + \epsilon_1 + \epsilon_2 & -\epsilon_3 \\ 0 & -\epsilon_2 & 1 + \epsilon_3 \end{pmatrix} \begin{pmatrix} P(DD) \\ P(BD + DB) \\ P(BB) \end{pmatrix}. \quad (8.5)$$

Eq. 8.5 can be used to remove the bias introduced by detection errors.

In our experiments, the detection calibration and correction is now handled automatically. A recurrent script fits experimental shot histograms, calculates the optimal thresholds  $(t_1, t_2)$  and the measurement errors  $(\epsilon_0, \epsilon_1, \epsilon_2, \epsilon_3)$  assuming Poissonian statistics. The measured probabilities  $P(DD), P(BD + DB), P(BB)$  are then processed using Eq. (8.5), and the resulting estimates  $P(00), P(01 + 10), P(11)$  are presented to the user. We typically set  $\epsilon_0 = \epsilon_1 = 0$ , since their effect is negligible in the Poissonian model. The MS gate analysis in Sec. 8.3 uses readout error estimates that account for ion decay. For the thresholds  $(t_1, t_2) = (32, 90)$ , which minimise the sum of all detection errors in Fig. 8.6 (bottom, dashed lines), these are  $(\epsilon_0, \epsilon_1, \epsilon_2, \epsilon_3) = (2, 0.1, 9, 10) \times 10^{-4}$ .

### 8.1.5. Axial stretch mode heating

In order to create high-fidelity entangled states despite anomalously high COM mode heating rates (Sec. 5.6.4), the MS gate is performed on the axial stretch mode, which is insensitive to electric-field noise that is common-mode to both ions [Kin98].

A general normal mode at frequency  $\omega_\alpha$  of a two-ion crystal experiences motional heating at a rate given by [Hom13]:

$$\dot{n}_\alpha = \frac{e^2}{4\hbar m \omega_\alpha} S_{E,\alpha}(\omega_\alpha)$$

where  $m$  is a mass of a single ion and

$$S_{E,\alpha}(\omega) = 2 \int_{-\infty}^{+\infty} dt \langle \xi(0) \xi(t) \rangle e^{i\omega t}$$

with

$$\xi(t) = \mathbf{E}(t) \cdot \mathbf{e}_\alpha$$

where  $\mathbf{E}(t) = (\vec{E}_1(t), \vec{E}_2(t))$  are the electric fields experienced ion 1 and 2 respectively, and  $\mathbf{e} = (\vec{e}_1, \vec{e}_2)$  are the eigenvectors of the normal mode in question. Now consider specifically the axial normal modes, whose eigenvectors only have a component along  $\hat{x}$ . We then find

$$\xi(t) = \mathbf{E}_x(t) \cdot \mathbf{e}_{x,\alpha}$$

where  $\mathbf{E}_x(t) = (E_{x,1}(t), E_{x,2}(t))$  are the x-components of the electric field at the ion 1 and 2 respectively,  $\mathbf{e}_{x,\text{com}} = (1, 1)/\sqrt{2}$  for the centre-of-mass mode and  $\mathbf{e}_{x,\text{str}} = (1, -1)/\sqrt{2}$  for the stretch mode. In the first-order expansion, we write  $\mathbf{E}_x(t) = (E_x(t), E_x(t) + \Delta x \partial E_x(t)/\partial x)$ , where  $\Delta x$  is the ion-ion spacing. Up to the leading order in  $\Delta x$ , the heating rates of the COM and stretch mode are then given by:

$$\begin{aligned} \dot{n}_{\text{com}} &= 2 \frac{e^2}{4m\hbar\omega_{\text{com}}} S_{E_x}(\omega_{\text{com}}) \\ \dot{n}_{\text{str}} &= \frac{1}{2} \frac{e^2}{4m\hbar\omega_{\text{str}}} S_{E_{xx}}(\omega_{\text{str}}) \Delta x^2, \end{aligned}$$

where

$$\begin{aligned} S_{E_x}(\omega) &= 2 \int_{-\infty}^{+\infty} dt \langle E_x(t) E_x(0) \rangle e^{i\omega t} \\ S_{E_{xx}}(\omega) &= 2 \int_{-\infty}^{+\infty} dt \left\langle \frac{\partial E_x(t)}{\partial x} \frac{\partial E_x(0)}{\partial x} \right\rangle e^{i\omega t}. \end{aligned}$$

Note that the two-ion COM heating rate  $\dot{n}_{\text{com}}$  is twice that of the single ion at the same frequency (Eq. (5.9)).

Consider an electric-field noise source distance  $d$  away from the ion. Dimensional analysis suggests that the ratio between the field and its gradient at the ion position [Win98]

$$\frac{\partial E_x(t)}{\partial x} = \alpha \frac{E_x(t)}{d},$$

where  $\alpha \sim 1$ . We might expect smaller values of  $\alpha$  for a field source with larger spatial extent along  $\hat{x}$ , and a larger  $\alpha$  for a source of size which is small compared to  $d$ . With this definition, we find that the ratio of the two-ion stretch and COM heating rates:

$$\frac{\dot{n}_{\text{str}}}{\dot{n}_{\text{com}}} = \frac{1}{4} \alpha^2 \frac{\Delta x^2}{d^2} \frac{\omega_{\text{str}}}{\omega_{\text{com}}} \frac{S_{E_x}(\omega_{\text{str}})}{S_{E_x}(\omega_{\text{com}})},$$

with  $\Delta x = (2e^2/4\pi\epsilon_0 m\omega_{\text{com}}^2)^{1/3}$ . As an order-of-magnitude estimate, consider the noise spectral density which falls off as  $1/f$ , such that  $S_{E_x}(\omega_{\text{com}})/S_{E_x}(\omega_{\text{str}}) = \sqrt{3}$ . An explicit calculation for electric dipoles homogeneously distributed across the trap surface yields  $\alpha^2 = 15/4$  [Tal15]. In that model, assuming values of  $\Delta x = 5\text{ }\mu\text{m}$  and  $d = 50\text{ }\mu\text{m}$  which approximates the conditions in our trap, we find an expected ratio of:

$$\frac{\dot{n}_{\text{str}}}{\dot{n}_{\text{com}}} \approx \frac{1}{300} \quad (8.6)$$

Thus indeed we expect a significant suppression of heating rates on the stretch mode.

We found it initially challenging to experimentally measure  $\dot{n}_{\text{str}}$ , since the large delay times necessary to resolve small heating rates resulted in significant mode frequency drifts (Sec. 8.2.2) and spectator mode heating. In the trap #4, where the MS gates described in Sec. 8.3 were performed, we measured

$$\dot{n}_{\text{str}} = 60(30) \text{ quanta/s} \quad (8.7)$$

at  $\omega_{\text{str}} = 2.2\text{ MHz}$  using the BSB fitting method (Sec. 5.6.1). Deteriorating two-ion lifetimes and loading rates prohibited us from constraining this quantity further. Given a single-ion heating rate  $\dot{n} \approx 2000 - 3000$  quanta/s at the same COM frequency (Sec. 5.6.4), this corresponds to  $\dot{n}_{\text{str}}/\dot{n}_{\text{com}} \approx 1/80$ , though with significant uncertainty. The slight suppression of the reduction compared to the order-of-magnitude estimate in Eq. (8.6) is likely related to the increased axial electric-field noise density at  $\omega > 1.5\text{ MHz}$ . Indeed, given the experimental observation that the single-ion heating rate is approximately independent of the centre-of-mass frequency (Sec. 5.6.4), we find that  $S_{E_x}(\omega_{\text{com}})/S_{E_x}(\omega_{\text{str}}) = 1/\sqrt{3}$ . This suggests a suppression ratio of  $\dot{n}_{\text{str}}/\dot{n}_{\text{com}} \approx 1/100$ , close to the experimental observation.


Longer ion lifetimes in the trap #5 allowed us to decrease the  $\dot{n}_{\text{str}}$  measurement uncertainty. The method of choice for those measurements was the sideband ratio test, which was performed by scanning the frequencies of a fixed-duration pulse near the BSB or the RSB resonance for a variable delay time. This allowed us to extract the sideband ratios despite continual drifts of motional frequencies. The expressions for the temperature are modified compared to the single-ion case in Sec. 5.6.1. For the system prepared in  $|11\rangle$ , we record the probabilities  $P_R(00)$ ,  $P_B(00)$  to find the ions in  $|00\rangle$  after a period of evolution under a drive resonant with the RSB or the BSB respectively. The temperature can be estimated from [Hom06]:

$$\bar{n} = \frac{R + \sqrt{R}}{1 - R} \quad \text{where} \quad R = \frac{P_R(00)}{P_B(00)}$$

The sideband ratio tests in the trap #5 indicated a stretch mode heating rate broadly consistent with the value obtained in the trap #4, although the results were not always reproducible. Worryingly, certain measurements hinted at a relation between  $\dot{n}_{\text{str}}$  and the duty cycle of the  $\lambda = 729\text{ nm}$  trap-integrated laser. One such observation is shown in Fig 8.7.



Each shot of the experiment, after the state readout, we add an additional  $\lambda = 729$  nm



**Figure 8.7:** Axial stretch mode heating rate measurement in the trap #5. The temperature is estimated with the sideband ratio method. The linear fits imply a heating gate in the range of 10 – 50 quanta per second, with a possible noise increase for increased average laser power in the waveguide. The increased mode temperature after EIT cooling for the data in red was likely an error related to trap potential drifts.

pulse with power equal to that used for an MS gate. Its length is set to either 0.1 ms or 3 ms (this was the method of the grating coupler charge stabilisation described later in Sec. 8.2.2). The measured heating rates:

$$\begin{aligned} \dot{n}_{\text{str}} &= 16(7) \text{ quanta/s} && 0.1 \text{ ms of } 729 \text{ per shot} \\ \dot{n}_{\text{str}} &= 43(5) \text{ quanta/s} && 3 \text{ ms of } 729 \text{ per shot} \end{aligned}$$


hint at a heating rate increase caused by increasing the integrated-waveguide laser power (the values in brackets indicate a  $1\sigma$  confidence interval of the fit). We repeated the measurement on the single-ion COM mode, where we consistently observed identical heating rates in both scenarios. Further experiments are necessary to establish whether differential electric-field noise is really affected by the laser pulses, although field drifts associated with laser duty cycle changes (Sec. 8.2.2) make such studies challenging. Indeed, the same data in Fig. 8.7 could be used to argue that increasing the laser power improves the mode temperature after EIT cooling, which is likely not the case but only an artefact of experimental drifts associated with such heating rate measurements.

## 8.2. System tuning and MS gate calibration

### 8.2.1. Substrate photoconductivity

The first single-ion experiments on the motional modes indicated significant transient frequency drifts following the turn-off of the cooling light. This is shown in Fig. 8.8 (red

points), where Rabi spectroscopy reveals a  $\approx 5$  kHz shift of the axial COM mode in the initial 10 ms following Doppler cooling and optical pumping. The immediate interpretation<sup>2</sup> of



**Figure 8.8:** Axial COM frequency as a function of wait time after Doppler cooling. (Red) Frequency drifts recorded initially (Blue) Frequency drifts disappears with 4 mW of light at  $\lambda = 780$  nm inserted into the waveguide loop.

this observation was that undesired scattering of the cooling beam into the silicon substrate causes electron-hole pair production. Owing to its large bandgap, silicon absorbs all the wavelengths used in our experiment. Following a pair production event, these charges diffuse through the substrate until they recombine over a timescale known as the free-carrier lifetime. The line-of-sight from the ion to the substrate at the grating couplers make us susceptible to its electric fields.

The solution we employed was to deliberately enhance this effect by permanently scattering photons into the substrate. This was achieved by inputting 4 mW of light from a laser diode at  $\lambda = 780$  nm directly into the waveguide loop (Sec. 4.6.4). We found this power to be sufficient to remove all observable frequency transients (Fig. 8.8, blue points). Our interpretation is that the scattered light causes sufficient substrate photoconductivity to effectively short it to the ground plane. In that case, any additional induced charges quickly flow to the ground plane, leaving the trap potential unaffected. This solution was in use for all the experiments in the trap #4. In the trap #5, we have not recorded comparable drifts (presumably due to improved laser alignment) and have since put the  $\lambda = 780$  nm source on stand-by<sup>3</sup>.


### 8.2.2. Light-induced charging

Further mode frequency drifts became observable upon increasing the  $\lambda = 729$  nm pulse powers to  $P_{\text{ion}} \gtrsim 1$  mW at the ion. At that level, turning on the qubit laser would cause a slow drift of the stray curvature at a timescale of seconds to minutes. This effect was

<sup>2</sup> Immediate to K. Mehta at least!

<sup>3</sup> In hindsight, this might have been premature, as  $\lambda = 729$  nm pulses themselves will probably induce the same effect, which would be less obviously noticeable

observable already with a single ion, and for all motional modes, as shown in Fig. 8.9.




**Figure 8.9:** Example observations of single-ion mode frequency drifts following a sudden turn on (red lines) or a sudden turn off (blue lines) of the trap-integrated light at  $\approx 1$  mW emerging from the coupler. (Top) Lower radial mode (Bottom left) Axial mode (Bottom right) Upper radial mode.

Qualitatively, we could observe this charging as a combination of two effects operating at different timescales. The first effect was a frequency jump upon the change of laser power. This would occur within 10 seconds after the change. This effect was often very pronounced when turning the laser on, and not so significant when turning the laser off, although systematic experiments were extremely challenging due to the long relaxation times of the system. Following the sudden jump, the frequency would proceed to drift over the next 1 – 5 minutes, before converging to a quasi-steady-state. If at that point the laser power is returned to the original value, the frequency would return closer to the original value as well, but never exactly, indicating significant stray charge hysteresis. Most commonly, increasing the laser power led to a long-term decrease of the trap frequencies (which we associated with positive coupler charging), although deviations from this behaviour could also be observed. The charging is definitely associated with the individual coupler, as passing high laser powers through couplers in other zones leaves the potentials unaffected.

The presence of two separate processes, as well as the overall drift timescale, is reminiscent of the behaviour observed in [Har10], which systematically studied the processes of laser charging of the trap surface and dielectrics. In that study, laser illumination at  $\lambda = 375$  nm and  $\lambda = 423$  nm lead to simultaneous negative charging (electron trapping) of trap electrodes,

and positive charging (electron ejection) of the dielectrics. The overall process was initially mainly driven by the negative charging of trap electrodes (at a timescale of  $2 - 16$  s), but eventually dominated by the positive charging of the dielectric (at a timescale of  $38$  s). In this study, like in other ones [Wan11; Van16; Ivo20], strong charging was only measured for blue and UV light. However, [Har10] did report an electrode (though not dielectric) charging with  $\lambda = 729$  nm incident light, which occurred at a rate  $\approx 100\times$  smaller than the one for blue and UV light at the same power. Our experiments indeed suggest that milliwatt-level  $\lambda = 729$  nm laser light moving through trap dielectrics leads to both negative and positive charging of the trap structures.

This issue highlights the need for dielectric shielding in future traps with integrated optics [Nif20]. In the meantime, we developed methods to partially mitigate the resulting frequency drifts. During the MS gate experiments in the trap #4, we combined the  $\lambda = 729$  nm pulses with always-on  $\lambda = 780$  nm light from a laser diode (the same as used to induce substrate photoconductivity, Sec. 8.2.1). This decreased the magnitude of frequency jumps, but the frequency still had to be calibrated every 15 seconds (Sec. 8.2.3). This performance was improved in the trap #5 by approximately stabilising the average power of  $\lambda = 729$  nm pulses per experiment. To that end, we add a milliwatt-level  $\lambda = 729$  nm laser pulse at the end of every experimental shot (after the qubit readout). The pulse duration is calculated such that every shot of every different experiment uses the same total high-power pulse length. For experiments with significantly different total shot length<sup>4</sup> we manually adjust the final pulse parameters to match the average optical power on a photodetector. The



**Figure 8.10:** Single-ion axial mode frequency drifts following a sudden turn on (red line) of the trap-integrated light at  $\approx 1$  mW emerging from the coupler. (Red) Before charge stabilisation, we observe motional frequency jumps as in Fig. 8.9. (Blue) With the charge stabilisation active, no motional frequency jumps can be observed.

resulting sequences exhibit no more observable motional frequency jumps, as shown in Fig. 8.10 by comparing the effect of starting an MS gate experiment with (blue) and without (red) the stabilisation pulse.

<sup>4</sup> These are typically experiments with long delays, such as Ramsey measurements of heating rate tests

### 8.2.3. Recurrent calibrations

In our setup, frequent drifts mean that a number of parameters have to be continuously adjusted before the MS gate can be calibrated. This is achieved through a Python script that keeps track of the time since the last parameter calibration. After that time is past, the current experiment is interrupted, a calibration experiment is executed and processed, and appropriate parameters are updated in the interrupted experiment. Occasional fit failure can occur if the parameter in question has drifted far enough from the previous value, or due to glitches in the experiment. In such a case, the calibration is automatically repeated with the same resolution but over a larger parameter range.

Finally, each automatic calibration experiment includes a number of dependencies [Mar20]. In the event of a few consecutive calibration failures in a row, the script pauses the calibration attempt and instead triggers a re-calibration of the dependencies. This approach is efficient in dealing with non-stationary noise, i.e. parameters whose variance changes over time. For example, it is wasteful to program a frequent check of the dipole optical pumping error, which usually stays unchanged over any given day. However, in the event of a suspected problem, a calibration experiment should be triggered as soon as possible.


### Frequency calibrations

Light-induced charging, discussed in Sec. 8.2.2, is by far the dominant source of experimentally relevant drifts. In the trap #4, the residual frequency drifts were mitigated by re-calibrating the axial frequency with a resolution of  $\approx 20$  Hz every  $\tau_{\text{cal}} = 15$  s. The histogram of calibrations obtained after the data-taking period can be fitted to obtain a mean drift of

$$\langle |\omega_{\text{str}}(t + \tau_{\text{cal}}) - \omega_{\text{str}}(t)| \rangle \approx 2\pi \times 200 \text{ Hz} \quad (8.8)$$

Besides eliminating the jumps associated with starting an experiment, improved coupler charging control in the trap #5 reduced this error, allowing us to decrease the time between calibrations to  $\tau_{\text{cal}} = 45$  s while keeping the noise variance approximately unchanged. The histogram of drifts in  $\omega_{\text{str}}$  obtained during a typical experimental day is shown in Fig. 8.11 (left). We have not yet established to what extent these residual fluctuations are exacerbated by laser-induced charging.

The next parameter which requires frequent recalibration is the optical qubit carrier frequency  $\omega_0$ . These drifts are dominated by slow changes in the external B-field, likely caused by temperature fluctuations of the permanent magnet “coils” (Sec. 3.4.3). While the drifts in  $\omega_0$  (Fig. 8.11, right) are smaller than the drifts of  $\omega_{\text{str}}$ , they can nonetheless introduce significant errors when left unchecked for a long time. Therefore, we periodically calibrate the bare optical qubit frequency through low-power Rabi spectroscopy with a



**Figure 8.11:** Histogram of frequency drifts between calibrations on a typical day in trap #5. (Left) Axial stretch mode frequency drifts, recalibration period of  $\tau_{\text{cal}} = 45$  s (Right) Optical qubit carrier frequency drifts, recalibration period of  $\tau_{\text{cal}} = 120$  s

resolution of  $\approx 15$  Hz. For a recalibration period of  $\tau_{\text{cal}} = 120$  s, we find a typical drift of

$$\langle |\omega_0(t + \tau_{\text{cal}}) - \omega_0(t)| \rangle \approx 2\pi \times 50 \text{ Hz}.$$

Very commonly, the histograms for both  $\omega_{\text{str}}$  and  $\omega_0$  cannot be well-fitted by a normal distribution [Sun19]. Averaging over a long data acquisition period, we typically find a normally-distributed centre peak surrounded by several outliers. These “black swan” events correspond to sudden jumps in the stray magnetic field or electric-field curvature and pose a serious challenge to experimental stability. So far, the only identifiable cause of sudden changes are nearby physicists, and their movements inside the laboratory (especially opening and closing the doors) can often be correlated with sudden magnetic-field changes. However, the majority of sudden frequency changes remain unexplained.

### Readout calibration

Slow drifts in  $\lambda = 397$  nm detection laser frequency and power (Sec. 4.3.3) lead to changes in detection errors. A recurrent calibration drives a long Rabi oscillation (typically every  $\tau_{\text{cal}} = 5$  min), processes the detection histograms and sets the optimal thresholds based on Eq. 2.3.6. Furthermore, it calculates the readout errors  $\epsilon_2$  and  $\epsilon_3$  from Eq. (8.4) and alerts the user if any one of them exceeds a hard-coded threshold. Finally, the Rabi oscillation trace is used to determine whether the Rabi frequency is still well-balanced (Sec. 8.1.3), or whether the ions’ position has drifted.

### Ions’ position drifts


In ion trap systems with free-space laser delivery, beam-pointing fluctuations are often suspected to cause significant Rabi frequency errors [Sch13; Bru19]. As discussed in

Sec. 5.5.2, the measured dipole Rabi frequency noise indicates a very low level of beam-pointing fluctuations in the integrated-waveguides trap, and the residual power drifts occur on timescales slower than individual experiments. However, two ions are a particularly sensitive probe of beam alignment, and the MS gate requires the differential Rabi frequency to be minimised<sup>5</sup>. This necessitates recurrent calibrations of the ions' position along the trap axis.

Once the approximate position  $x = x_0$  and average Rabi frequency  $\Omega$  are found as described in Sec. 8.1.3, we perform a faster experiment to compensate for small position drifts. Two ions are prepared in  $|11\rangle$  and evolve under a carrier drive for a period of

$$t \approx 11.75 \times \pi / \Omega.$$

for a range of values of  $x$ . A typical result of a subsequent population measurement is shown in Fig. 8.12. By choosing  $t$  slightly below an even multiple of the pi-time, the balanced



**Figure 8.12:** Fast calibration of ions' position along the trap axis. The centre of parabolic fit (dashed lines) at  $x_0 = 3.41 \mu\text{m}$  indicates a position where the Rabi frequency imbalance (Eq. (8.2)) is minimised.

position  $x_0$  can be found as the value of  $x$  with the largest population in  $P(11)$  and the lowest population in  $P(01 + 10)$  (the dark-state population  $P(00)$  is typically negligible), regardless of small errors in  $\Omega$ . We typically scan  $x$  over a span of  $\approx 500 \text{ nm}$ , and 20 data points with 200 experiments per point are sufficient to position the ions with a precision of  $\approx 10 \text{ nm}$ . This calibration takes approximately a minute, limited by long trap voltage source response time.

Unlike the motional and carrier frequencies, the drifts in  $x_0$  are predominantly non-stationary. The typical drift pattern proceeds as follows. Following ion loading, the value of  $x_0$  drifts significantly, on the order of  $\approx 500 \text{ nm}$  over  $\approx 5 \text{ min}$ . During that period, we do not attempt entangling experiments, and instead, calibrate all the necessary parameters

<sup>5</sup> This is because differential Rabi frequency causes a differential AC Stark shift, which then cannot be perfectly compensated by a global frequency shift

while monitoring  $x_0$ . At some point,  $x_0$  reaches an approximate steady-state, with typical drifts of  $< 50$  nm over 10 min. We then begin the experiments, and correct for small drifts in  $x_0$  only every  $\approx 30$  min or if the readout error calibration reports a significant Rabi frequency imbalance. Note that every update of  $x_0$  needs to be followed by a calibration of  $\omega_{\text{str}}$ , which has a small correlation with  $x_0$  due to trap modelling errors. However, the MS gate error is dominated by sudden jumps in  $x_0$ , which occur infrequently but lead to sudden infidelity spikes. I hope that the recent upgrade to fast voltage sources (Sec. 3.4.4) can improve this error by speeding up the position calibrations.

#### 8.2.4. MS gate calibration

As discussed in Sec. 2.4, an MS pulse involves simultaneous application of two drives of equal Rabi frequencies, near-detuned from the RSB and the BSB of the gate mode. The pulse is shaped over  $\approx 2.5$   $\mu\text{s}$  by smoothly ramping the power sent to the double-pass AOM, while the bichromatic drive is achieved by adding two RF tones on a single fibre-coupled AOM (Sec. 4.4.7). The first calibration step is to equalise the RSB and BSB Rabi frequencies, which can be done by matching their powers on a photodetector. The pulse powers can be either measured and adjusted individually or through a beat-note measurement with both drives switched on simultaneously. Following that calibration, we only change the total pulse power, and always assume that the individual tone powers remain equal to each other.

The remaining calibrations are performed using the ions while the recurrent calibrations are running in the background. The ions are cooled, prepared in  $|11\rangle$ , and we monitor the population at a desired gate time  $t_{\text{MS}}$ . We set the BSB drive frequency to  $\omega_0 + \omega_{\text{str}} + \delta + \delta_{\text{com}}$  and the RSB drive frequency to  $\omega_0 + \omega_{\text{str}} - \delta + \delta_{\text{com}}$ . The common-mode detuning  $\delta_{\text{com}}$  is introduced to compensate for the AC Stark shift of the carrier transition caused by the MS drive. Since the BSB and RSB drives are close to symmetrically detuned from  $\omega_0$ , their quadrupole AC Stark shifts are mostly cancelled, and the remaining shift is predominantly caused by coupling to dipole-allowed transitions [Kir09a; Häf03]. For a given MS pulse time  $t_{\text{MS}}$ , we find a combination of  $\delta$  and  $\delta_{\text{com}}$  which minimises the odd population  $P(01 + 10)$  at time  $t_{\text{MS}}$ . The obtained value of  $\delta$  typically deviates slightly from the theoretical value of  $\delta = 2\pi/t_{\text{MS}}$  due to a non-rectangular pulse shape.

In the next calibration step, the Rabi frequency  $\Omega$  is adjusted to the optimal value of  $\Omega_0 \approx \delta/2\eta$ , where  $\eta$  is the Lamb-Dicke parameter for the gate mode (Eq. 2.4.2). The signature of the correct Rabi frequency calibration is that  $P(11) = P(00)$  at time  $t_{\text{MS}}$ , which is a pre-requisite for a maximally entangled state. However, changes of  $\Omega$  require simultaneous adjustments of  $\delta_{\text{com}} \propto \Omega^2$  to compensate for varying AC Stark shifts. In the event that  $P(11) > P(00)$  at time  $t_{\text{MS}}$  for all available Rabi frequencies, the laser power is insufficient to drive an MS gate of duration  $t_{\text{MS}}$ , and the full calibration is repeated at a longer gate time  $t'_{\text{MS}} > t_{\text{MS}}$ . On the other hand, if  $\Omega$  can be comfortably adjusted to obtain  $P(11) < P(00)$  at time  $t_{\text{MS}}$ , it is beneficial to decrease the total gate time to  $t'_{\text{MS}} < t_{\text{MS}}$ , as shorter gates result in lower errors (Sec. 8.3.2). The calibration procedure is therefore




iterated until a gate of optimal length (given the available  $\Omega$ ) is achieved.

In the trap #4, we found the following parameters for an MS gate. At a stretch mode frequency of  $\omega_{\text{str}} = 2\pi \times 2.03$  MHz, we drive a maximally entangling operation in  $t_{\text{MS}} = 66 \mu\text{s}$  with

$$\begin{aligned} \delta &= 2\pi \times 16.7 \text{ kHz} && \text{sideband detuning,} \\ \delta_{\text{com}} &= 2\pi \times 18.3 \text{ kHz} && \text{AC Stark shift.} \end{aligned}$$

Given the Lamb-Dicke parameter  $\eta = 0.028$ , this corresponds to an average carrier Rabi frequency  $\Omega_0 \approx 2\pi \times 270$  kHz.




**Figure 8.13:** Fast MS gate recalibration scan. The AOM amplitude is adjusted, varying the overall Rabi frequency. The linear fits to  $P(11)$  and  $P(00)$  (dashed lines) intersect at the AOM amplitude corresponding to a Rabi frequency  $\Omega = \Omega_0$ . The parabolic fit to  $P(01 + 10)$  (dashed parabola) shows that the odd population is minimised as the same AOM amplitude.

While the initial MS gate calibration is time-consuming and iterative, subsequent calibrations are significantly simpler. This is because, even though parameter drifts originate from multiple different processes, they all affect the MS interaction in only one of two ways. First, drifts in light polarization, the ions' position and micromotion modulation all affect the carrier Rabi frequency  $\Omega$ , changing the ideal gate time and the total AC Stark shift. Second, drifts of trap potential curvature lead to drifts in  $\omega_{\text{str}}$ , affecting the Lamb-Dicke parameter  $\eta$ , and in turn the ideal gate time.

This observation suggests a simple method of gate re-calibration. The MS pulse time  $t_{\text{MS}}$  and detunings ( $\delta, \delta_{\text{com}}$ ) are held constant. First,  $\omega_{\text{str}}$  is set to the value where the gate was initially calibrated by adjusting the trapping potentials. This does not need to be done very precisely, as  $\eta$  is only weakly dependent on  $\omega_{\text{str}}$ . Second, we scan the value of  $\Omega$  by varying the overall laser power. We find that  $\Omega = \Omega_0$ , where  $P(00) = P(11)$ , is then also the Rabi frequency where the AC Stark shift is compensated, and therefore  $P(01 + 10)$  is minimised. This is shown in Fig. 8.13. Using this method, we typically only require a single scan to recalibrate an MS gate..

### 8.3. MS gates

The MS gate experiment combines all the elements described in this chapter thus far. Two ions are trapped and ground-state cooled as discussed in Sec. 8.1.2 and optically pumped into  $|11\rangle$  on dipole and quadrupole transitions (Sec. 5.2). The calibrated MS pulse (Sec. 8.2.4) is then applied with recurrent calibrations (Sec. 8.2.3) running in the background. The two-ion state is measured as described below in Sec. 8.1.4.




**Figure 8.14:** Measured two-ion populations as a function of the MS drive duration. Dashed lines represent theoretical evolution for an MS drive Hamiltonian with axial stretch mode cooled to  $\bar{n} = 0.03$ . Blue vertical line at  $t_{\text{MS}} = 66 \mu\text{s}$  shows the time where  $P(00) = P(11)$ , and the ions' state is closest to maximally entangled.

Fig. 8.14 shows the measured two-ion state as a function of the MS pulse time  $t$ . The MS gate time  $t_{\text{MS}} = 66 \mu\text{s}$  is highlighted with a blue line. Dashed black lines are the theoretical traces for an ideal MS gate (Sec. 2.4), obtained by solving the Lindblad master equation (Eq. (2.17)) for the Hamiltonian (Eq. (2.25)) with the axial stretch mode in a thermal state of  $\bar{n} = 0.03$ . We see that the experimental values are in excellent agreement with the theoretical expectation, not just at the gate time but throughout the pulse. Slight systematic offsets are likely caused by the imperfections described in Sec. 8.3.2, as well as by pulse shaping which was not included in the simulation.

#### 8.3.1. Measured performance

We quantify the performance by measuring the Bell-state fidelity using Eq. (2.35) at  $t_{\text{MS}}$ . The measurements of the odd population  $P(01 + 10)$  and the parity following a  $\pi/2$  analysis pulse  $\langle \sigma_\phi \sigma_\phi \rangle$  are shown in Fig. 8.15.

The odd population measurement can be analysed in two ways. First, a parabolic unweighted least-squares fit (solid lines) around the gate time gives an offset from the origin. Alternatively, we can average over the middle five data points, whose mean should



**Figure 8.15:** Bell-state fidelity measurement. (Left) Odd population  $P(01 + 10)$  close to the gate time of  $t_{\text{MS}} = 66 \mu\text{s}$ , and a parabolic fit (solid line) (Right) Parity  $\langle \sigma_\phi \sigma_\phi \rangle$  measurement following a  $\pi/2$ -pulse via a variable phase  $\phi$ . Solid line shows a fit of  $\langle \sigma_\phi \sigma_\phi \rangle = C \cos(2\phi - \phi_0)$  (Eq. (2.36)).

not exceed the true minimum value of  $P(01 + 10)$ . Both methods agree within error bars, and we used the latter as the more conservative estimate to avoid a possible bias from the least-squares fit (see below). This allows us to obtain the even population at the gate time  $t_{\text{MS}} = 66 \mu\text{s}$  of:

$$P(00) + P(11) = 0.9946(10), \quad (8.9)$$

where the quantity in brackets is the  $1\sigma$  confidence interval obtained by averaging five points whose individual uncertainties are given by the quantum projection noise (Sec. 2.3.6).

The parity contrast is extracted with a maximum-likelihood estimation (MLE) fit to the dataset shown in Fig. 8.15 (right). The MLE method provides an unbiased estimator of the true contrast [Bal14], which is not the case for the least-squares method due to non-normal statistics<sup>6</sup> of the data points close to  $\langle \sigma_\phi \sigma_\phi \rangle = \pm 1$ . Instead, the parity data is analysed based on the following principle. For every analysis pulse setting  $\phi_i$ , the experiment is repeated  $N$  times. We classify the outcomes into two classes: “success” is when the parity was found to be even, i.e. ions’ state was  $|00\rangle$  or  $|11\rangle$ , and “failure” when the parity was odd, i.e. the state was  $|01\rangle$  or  $|10\rangle$ . This way, we can consider every experimental shot as an independent Bernoulli trial [Hug10]. For every setting, we assume the success probability to be given by

$$p_i(C, \phi_0) = \frac{1}{2}[1 + C \cos(2\phi_i - \phi_0)],$$

where  $C$  and  $\phi_0$  are the parameters to be estimated. This is expected to be exact for any two-qubit pure or mixed state.

The MLE analysis is performed as follows. For every setting  $\phi_i$ ,  $i = 1, 2, \dots, m$  we extract

<sup>6</sup> In fact, the least-squares method is derived from the maximum-likelihood formalism under the assumption of normally-distributed points [Hug10]


the number of successful trial outcomes  $k_i$  and calculate the corresponding probability from the probability density function of a binomial distribution:

$$P_i(C, \phi_0) = \binom{N}{k_i} [p_i(C, \phi_0)]^{k_i} [1 - p_i(C, \phi_0)]^{N-k_i}.$$

We then calculate the total log-likelihood of the experimental outcomes given the model parameters

$$\log \mathcal{L}(C, \phi_0) = \sum_{i=1}^m P_i(C, \phi_0).$$

The MLE is completed by finding<sup>7</sup> a set of parameters  $C \in [0, 1]$  and  $\phi_0 \in [-\pi, \pi]$  that maximises  $\log \mathcal{L}(C, \phi_0)$ . To estimate the fit error, we repeat the MLE analysis on a number of bootstrap samples [Hug11]. Each bootstrap set consists of a set of  $m$  probabilities  $P'_i(C, \phi_0)$ ,  $i = 1, 2, \dots, m$  sampled with replacement from the original probabilities  $P_i(C, \phi_0)$ . The



**Figure 8.16:** Bootstrapped maximum-likelihood estimates of the parity curve contrast. Solid line shows the mean estimated contrast, and dashed lines indicate  $\pm 1$  standard deviation in the estimate.

MLE of  $C$  for 2000 bootstrap samples are shown in Fig. 8.16. The contrast is estimated to be:

$$C = 0.9922(23) \tag{8.10}$$

where the error is the standard deviation of the bootstrap samples (dashed lines). By comparison, a least-squares analysis returns  $C = 0.994(6)$ .

Combining the even population estimate in Eq. (8.9) with the parity contrast estimate in Eq. (8.10), we obtain the fidelity estimate (Eq. (2.35)):

$$F = 0.9934(13). \quad \text{Estimated fidelity.}$$

<sup>7</sup> The results below were obtained with the interior-point method, which is the default method for constrained local optimisation in Wolfram Mathematica

The data used for fidelity estimation was not corrected for readout errors. The reason is that a naive application of Eq. (8.5) corrects the readout probabilities to values which may be inconsistent with a binomial trial<sup>8</sup>, confusing the MLE. Instead, as a simpler estimate, we apply the correction matrix to the fitted parity curve, assuming  $P(00) = P(11)$  at all times. This increases the even population estimate at the gate time by  $6 \times 10^{-4}$  and the parity contrast by approximately  $14 \times 10^{-4}$ . The final fidelity estimate reads:

$$F = 0.9945(13) \quad \text{Readout-corrected fidelity.}$$

This is still a lower bound on the entangled state fidelity since it is not corrected for the state preparation error and the analysis pulse error. The former is expected to be at  $\approx 1 \times 10^{-4}$ , approximately twice that for a single ion. The latter only affects the contrast measurement, where it creates an error comparable with the contrast loss of a single-ion pi-pulse. The total contrast loss from expected from thermal Rabi-frequency fluctuations is  $\approx 4 \times 10^{-4}$ , but could be as high as  $10 \times 10^{-4}$  given the anomalously low depth of pi-pulses measured with a single ion (Sec. 5.5). Therefore, the analysis pulse creates a fidelity measurement error of  $\approx 2 - 5 \times 10^{-4}$ .

### 8.3.2. Error model

We now turn our attention to understanding the causes of the measured Bell-state infidelity. Following [Bal14], we tabulate the main expected error sources and their contributions to the Bell-state infidelity. We include errors that are expected to individually cause a fidelity loss of  $> 10^{-4}$ , and we round all the estimates to the nearest multiple of  $10^{-4}$ . For simplicity, error bars are only stated for the error sources with the largest uncertainties ( $\gtrsim 10^{-3}$ ).

#### Ion temperature and Rabi frequency fluctuations

While the MS drive fidelity is nominally independent of the gate mode temperature, this effect is only exact in the Lamb-Dicke approximation. Furthermore, warm spectator modes contribute to Rabi frequency fluctuations (Eq. (2.38)), which induce gate errors. Each shot, the gate error due to a Rabi frequency set to  $\Omega = \Omega_0 + \delta\Omega$  is given by:

$$\epsilon_\Omega = \frac{\pi^2}{4} \delta\Omega^2 / \Omega_0^2$$

where  $\Omega_0$  is the ideal Rabi frequency [Bal14]. This expression allows us to model the effect of both thermal and other slow-varying Rabi frequency fluctuations by setting  $\delta\Omega^2$  to the

---

<sup>8</sup> For any given setting  $\phi_i$ , the number of successful outcomes  $k_i = N \times p_i$ , where  $p_i$  is the probability of a successful outcome. Following an application of Eq. (8.5),  $k_i$  may not be given by a whole number, or we may even find  $k_i > N$ . A possible solution is to probabilistically correct individual shots. However, at that point, a Bayesian procedure becomes even more suitable.

noise variance  $\sigma_\Omega^2$ . The variance of shot-to-shot Rabi frequency fluctuations was measured in Eq. (8.3) by fitting the decay rate of carrier oscillations. This allows us to estimate the associated error as:

$$\langle \epsilon_\Omega \rangle = \frac{\pi^2}{4} \left( \frac{\sigma_\Omega}{\langle \Omega \rangle} \right)^2 = 3 \times 10^{-4}.$$

### Heating rates

Aside from increasing the Rabi frequency fluctuations, heating of the gate mode directly affects the gate performance. Modelled as shown in Tab. 2.3, the error caused by axial stretch mode heating rate  $\dot{n}_{\text{str}}$  is given by [Bal14]:

$$\epsilon_h = \frac{\dot{n}_{\text{str}} t_{\text{MS}}}{2}.$$

For  $t_{\text{MS}} = 66 \mu\text{s}$  and  $\dot{n}_{2,\text{str}} = 60(30)$  quanta/s determined in the trap #4 (Eq. (8.7)), we find the associated error of

$$\epsilon_h = 2(1) \times 10^{-3}.$$

This is likely the largest error source of the gate.

### Mode frequency drifts and cross-Kerr coupling

A constant offset  $\delta$  in the motional frequency creates an MS gate error given by:

$$\epsilon_\delta = \left( \frac{3}{16} + \frac{\bar{n}}{4} \right) \delta^2 t_{\text{MS}}^2 \quad (8.11)$$

where  $\bar{n}$  is the gate mode temperature [Bal14]. This expression allows us to model the effect of any slow-varying motional frequency noise, with  $\delta^2$  replaced by the noise variance  $\langle \delta^2 \rangle$ . We identify two dominant sources of detuning noise.

The first is due to slow drifts in motional frequencies. We model the continual frequency drifts between calibrations as a random walk of the gate detuning  $\delta$ . We assume  $\delta = 0$  immediately following a calibration at time  $t$ , and that the subsequent calibration at time  $t + t_0$  finds the frequency normally distributed with standard deviation  $\sigma_0$ . Then, the average standard deviation in detuning between calibrations [Rie03, Tab. 3.1]:

$$\sqrt{\langle \delta^2 \rangle} = \frac{1}{t_0} \int_t^{t+t_0} \sigma_0 \sqrt{\frac{t}{t_0}} = \frac{2}{3} \sigma_0$$

Using  $\sigma_0 = 2\pi \times 250 \text{ Hz}$  for a typical calibration histogram<sup>9</sup> (Eq. (8.8)),  $t_{\text{MS}} = 66 \mu\text{s}$  and

---

<sup>9</sup> Note that we converted from the average absolute deviation to standard deviation as described in [Gea35]

$\bar{n} = 0.03$  we find

$$\epsilon_\delta = 1.4(10) \times 10^{-3}.$$

where the error in brackets is a rough guess of typical variations of  $\sigma_0$  ( $\approx \pm 2\pi \times 100$  Hz), as well as an admission of deficiencies of the random-walk model.

The second source of detuning noise is the cross-Kerr coupling between the axial stretch mode and the radial rocking modes (Sec. 2.5.1). Using Eq. (2.39) with  $\omega_s = 2\pi \times 2$  MHz,  $\omega_{r1} = 2\pi \times 3.7$  MHz and  $\omega_{r2} = 2\pi \times 5.7$  MHz, we find the cross-Kerr coupling coefficients of  $\chi_1 = -2\pi \times 5.9$  Hz and  $\chi_2 = -2\pi \times 3.7$  Hz to the lower and the upper radial modes respectively. At the time of the data acquisition, Doppler cooling was not optimised as well as in Sec. 5.6.2, and we measured a temperature of  $\bar{n}_{r1} \approx 12$  for the lower radial mode. This creates axial stretch-mode frequency noise with a standard deviation of (Eq. (2.40)):

$$\sqrt{\langle \delta^2 \rangle} = 2\pi \times 100 \text{ Hz}$$

and Eq. (8.11) implies the corresponding error is given by:

$$\epsilon_{\chi_1} = 4 \times 10^{-4}.$$

The upper radial mode has a lower thermal occupancy of  $\bar{n}_{r2} \approx 5$  and is thus expected to contribute a negligible error of  $\epsilon_{\chi_2} = 3 \times 10^{-5}$ .

### Laser frequency noise


As discussed in Sec. 5.4.3 (Fig. 5.16), the dephasing at  $t_{\text{MS}}$  is predominantly caused by monochromatic noise with frequency  $\omega_0 = 2\pi \times 175$  Hz and amplitude  $A \approx 2\pi \times 185$  Hz. Since  $t_{\text{MS}} \ll 2\pi/\omega_0$ , we assume the laser detuning to be constant during any single experimental shot. I am not aware of any analytical expression for the gate error caused by a static detuning offset  $\delta_{\text{com}}$ , so I evaluate it numerically. The result is shown in Fig. 8.17.

Before evaluating the error, note that for any  $\delta_{\text{com}}$  the infidelity with respect to the target Bell state significantly exceeds the error with respect to some Bell state. This can be understood as follows: for  $\delta \neq 0$ , not only is the final state not maximally entangled (orange points are above zero), but also the relative phase of  $|00\rangle$  and  $|11\rangle$  changes with  $\delta$  (pushing the blue points even higher). The error associated with fluctuating detuning can be found through a parabolic fit (black line) to the error with respect to the target Bell state, resulting in an empirical formula for  $\bar{n} = 0$ :

$$\epsilon_{\text{com}} \approx 0.4\delta_{\text{com}}^2 t_{\text{MS}}^2. \quad (8.12)$$

For sinusoidal noise of amplitude  $A$  we find a noise variance of  $\langle \delta_{\text{com}}^2 \rangle = A^2/2$ . Combining

## 8. ENTANGLEMENT WITH INTEGRATED WAVEGUIDES



**Figure 8.17:** Simulated Bell-state error due to a common-mode frequency offset  $\delta_{\text{com}}$ , together with parabolic fits in Eq. (8.12). (Left) Error for a fixed gate time  $t_{\text{MS}} = 66 \mu\text{s}$  and variable  $\delta_{\text{com}}$  (Right) Error for a fixed detuning  $\delta_{\text{com}} = 2\pi \times 100 \text{ Hz}$  and a variable gate time  $t_{\text{MS}}$ .

these expressions together, we estimate the error associated with laser frequency noise as:

$$\epsilon_{\text{com}} \approx 1.2 \times 10^{-3}.$$

A notable complication is that the analysis pulse used to measure the fidelity is also applied at the detuning  $\delta_{\text{com}}$ , cancelling some of the apparent phase error. However, simulations suggest this to be a small correction ( $\approx 10\%$ ).

As discussed in Sec. 5.4.3, the remaining laser frequency noise may be low- or high-bandwidth. While the low-bandwidth noise would only cause negligible error at  $t_{\text{MS}} = 66 \mu\text{s}$ , high-bandwidth noise requires careful consideration. Ramsey measurements in Sec. 5.4.3 constrained the timescale of Markovian optical qubit dephasing to  $T_2 < 50 \text{ ms}$ . Setting an approximate lower-bound of the laser linewidth to  $\Delta\omega \gtrsim 2\pi \times 2 \text{ Hz}$  gives  $T_2 \lesssim 330 \text{ ms}$ . Simulations indicate that the associated error is then given by the range of

$$\epsilon_{T_2} \approx 2 \frac{t_{\text{MS}}}{T_2} \quad (8.13)$$

which is in the range of

$$4 \times 10^{-4} < \epsilon_{T_2} < 26 \times 10^{-4}$$

### Neglected noise sources

Finite optical qubit lifetime  $T_1 = 1.1 \text{ s}$  (Sec. 5.4.1) can be simulated to cause an error or  $\epsilon_{T_1} \approx 0.5 t_{\text{MS}} / T_1 \approx 3 \times 10^{-5}$ . We did not estimate the rate of motional dephasing, which contribute additional error. Residual Rabi frequency imbalance  $|\epsilon| \approx 10^{-2}$  leads to a small differential AC Stark shift between the ions of  $\approx 2\epsilon\delta_{\text{com}} = 2\pi \times 366 \text{ Hz}$ . This causes an error of  $7 \times 10^{-5}$  in the simulations. While simulations with rectangular pulses indicate



<b>Bell-state error source</b>	<b>Infidelity <math>\times 10^{-4}</math></b>
Gate mode heating	10 – 30
Motional frequency drifts	4 – 24
Laser frequency noise (Markovian)	4 – 26
Laser frequency noise (monochromatic)	12
Cross-Kerr coupling	4
Rabi frequency fluctuations	3
Others	1
<b>Total gate error sources</b>	<b>38 – 100</b>
<b>Measurement error source</b>	
State preparation	1 – 2
Analysis $\pi/2$ pulse	2 – 5
<b>Total measurement error sources</b>	<b>3 – 7</b>
<b>Measured error</b>	<b>42 – 68</b>

**Table 8.1:** Summary of expected Bell-state error sources and their contributions to the measured error. The total of the gate error sources is the sum of the lower- and upper- estimated limits, and not the quadrature sum.

a significant error due to off-resonant carrier excitations, shaping the rectangular pulses with an envelope of  $t = 2.5 \mu\text{s}$  reduces the simulated error to  $< 10^{-5}$ . Pulse shaping leads to a time-dependent AC Stark shift during the pulse turn-on and turn-off. In simulations, the associated error could be reduced to  $< 10^{-4}$  by adjusting the pulse time (I have not studied it carefully).

### Error summary

Table 8.1 summarises the individual error contributions. The error sources discussed in this section are sufficient to account for the observed fidelities. Note that the sum of the individual errors is only an order-of-magnitude estimate of their combined contribution. The biggest sources of uncertainty were the gate mode heating rate (which was not precisely measured), motional frequency drifts (which vary in magnitude over time) and the Markovian laser frequency noise.

### 8.3.3. Performance in the trap #5

While we haven't explicitly aimed at characterising and improving the MS gates in the subsequent trap, we used the MS interaction for the experiments described in Chapter 9.

Therefore, the measurements from the trap #5 can be used to address some of the uncertainties of the error budget. However, any comparisons between chips are necessarily approximate due to a number of systematic differences in the trap performance and in the experiment calibration. In addition strong effective trap potential anharmonicity observed only in trap #5 (Sec. 9.3.3), a reduced fibre-waveguide coupling led to an increased gate time of  $t_{\text{MS}} = 77 \mu\text{s}$ . The sideband cooling of the axial stretch mode (now at  $\omega_{\text{str}} = 2\pi \times 2.2 - 2.4 \text{ MHz}$ ) was working less well, with  $\bar{n} \approx 0.08$ , but the lower radial mode could be cooled to a lower Doppler temperature of  $\bar{n} \approx 6.5$  (Sec. 5.6.2)

### Phase-modulated gates

We anticipated that the use of phase-modulated gates (Sec. 2.4.3) with  $t'_{\text{MS}} = \sqrt{2}t_{\text{MS}} = 109 \mu\text{s}$  would lead to significant improvements in Bell-state fidelities due to suppression of important error sources in Tab. 8.1. These are:

1. Gate mode heating. For  $\dot{n} = 16 - 43$  quanta/s (Sec. 8.1.5) we expect a heating rate error  $\epsilon_h = \dot{n}t'_{\text{MS}}/4 = (4 - 12) \times 10^{-4}$ , a factor of  $\sqrt{2}$  suppression compared to a single-loop gate [Bal14].
2. Motional frequency drifts. For the same magnitude of slow frequency drifts, simulations suggested that phase modulation reduces the residual error to  $\epsilon_\delta = (1 - 5) \times 10^{-4}$ , approximately a factor 5 suppression compared to a single-loop gate [Hay11].
3. Cross-Kerr coupling. This error was expected to be negligible, both due to the phase modulation (which suppresses the effect of mode frequency variations) and the reduced lower radial mode temperature (which reduces mode frequency variations).

However, a prolonged gate duration may accentuate the error caused by the laser frequency fluctuations:

1. The monochromatic component of the laser frequency noise was simulated to cause an error of  $\epsilon_{\text{com}} = 16 \times 10^{-4}$ , approximately the same as for a single-loop gate.
2. The Markovian component of the laser frequency noise was expected to increase in error to  $\epsilon_{T_2} = (7 - 44) \times 10^{-4}$ , compared to  $\epsilon_{T_2} = (5 - 31) \times 10^{-4}$  for a single-loop gate.


The overall error was therefore expected in the range of  $(4 - 8) \times 10^{-3}$  for a two-loop phase-modulated gate, comparable to the expected single-loop error range of  $(3 - 9) \times 10^{-3}$ .

Measurements in the trap #5 revealed typical infidelities of  $\approx 10 \times 10^{-3}$  for both single-loop and phase-modulated two-loop gates. These results hint that the latter may be dominated by Markovian laser phase noise. Nonetheless, it is not clear why the error in the trap #5 was consistently at the upper limit of the expected noise, while in the trap #4 it was at the close to the lower limit. It could indicate a systematic change in the

laser spectrum or a novel error source. While phase-modulated gates did not improve the Bell-state fidelities, we found them beneficial for robustness to noise in long pulse sequences, and their experimental use is described in detail in Chapter 9

### Stability

During the data acquisition period in trap #4, we were not able to perform repeated measurements to verify the true variance of the fidelity. By contrast, MS gates were frequently performed in the trap #5 over six months. The Bell-state fidelity measurements, whose results are shown in Fig. 8.18, were taken with lower precision, usually only after a fast calibration, and for simplicity were analysed with least-squares analysis only. Nonetheless, the observed variations provide important data points regarding the gate stability.



**Figure 8.18:** Bell-state infidelities recorded for single-loop MS gates in the trap #5. Error estimates were obtained with least-squares fits to the odd population curves and the parity curves (as such in Fig. 8.15), and the error bars represent a  $1\sigma$  confidence interval. Yellow circle shows the second-class results, and blue circle highlights the third-class results (see main text)

Qualitatively, we find the results fall into three classes. In the first class, the results (uncircled) are consistent with  $F \gtrsim 99\%$  Bell-state fidelity<sup>10</sup>. In the second class (yellow circle), the data falls firmly in the range of 1 – 2% error. I speculate that it could be caused by variations in laser linewidth (such as reported in Sec. 5.4.3), which may suddenly change due to laser lock adjustments, and then stay unchanged for weeks. In December 2020, we observed a further fidelity deterioration to  $F < 98\%$  (blue circles). After extensive experimentation, we found it to be only fixed by changing the axial stretch mode frequency  $\omega_{\text{str}}$ . We speculate that the error increase was caused by the MS drive off-resonantly exciting a second-order spectator radial sideband. As the trapping potentials drift over weeks, we would infrequently observe similar fidelity drops. These could always be fixed by adjusting  $\omega_{\text{str}}$  over a few tens of kHz.

<sup>10</sup> Imagine your university required a score of  $\gtrsim 99\%$  for a first-class degree!



## 9. Optical pumping into a maximally entangled state

Pioneering theoretical work on quantum computing was partly motivated by a broader quest for *reversible computation* [Fey96]. Long before quantum computing, physicists became aware that indefinite improvements to computing speed and processor size may be limited by the increased energy use and the resulting heat dissipation. This is because chip components such as transistors are macroscopic objects, whose behaviour is fundamentally irreversible. This necessarily leads to heat dissipation, as formalised by Landauer’s principle [Lan61; Ben02]. On the other hand, atomic manipulation is in principle unitary and reversible. This means there is no lower bound on the amount of energy necessary to perform quantum computation [Ben73]. The state preparation and measurement steps, while irreversible, only add a constant factor of energy, which does not increase with increased computation length<sup>1</sup>.

However, this picture is incomplete due to the presence of unavoidable errors and noise, which cause non-unitary (or *dissipative*) dynamics. In the language of thermodynamics, this can be considered as a coupling between a “hot” environment and a quantum register [Ste03]. Left alone, the register will gradually “thermalise” with the environment, leading to increased entropy and information loss. In order to stabilise the “temperature” of the register, it is necessary to create an additional dissipative process which causes “heat flow” from the register to the environment. Thus, dissipation during computation becomes necessary to sustain quantum information. The most prominent example of this paradigm is QEC. There, the entropy of the quantum register is stabilised by coupling it to a “cold” ancilla. At the same time, ancilla preparation and stabiliser measurements cause a constant “heat flow” from the ancilla to the environment. Thus, in the circuit model, universal error-corrected quantum computation can be performed by interleaved application of unitary and non-unitary operations.

An alternative framework known as “dissipation engineering” allows for universal quantum computation without any unitary operations [Ver09]. In particular, controllable dissipation can be used to robustly prepare ground-states of different Hamiltonians, thus implementing adiabatic quantum computation. However, while adiabatic quantum computation is known to be universal, it is currently unknown whether it can be fault-tolerant [Aha04]. Nonetheless,

---

<sup>1</sup> In theory! In practice, present-day quantum computers are a significant energy drain. While modern computer CPU uses under 100 W for a clock speed of 3 GHz, a small trapped-ion processor uses around 7 kW of power at a two-qubit gate clock speed of 5 kHz [Pog21]

dissipation engineering is a promising framework for preparing specific desired quantum states [Poy96; Kra08; Ple99; Ver09; Tic14]. Within this model, useful quantum states are created in the steady-state, making the process more robust to imperfections [Kas11; Mor15], autonomous, and resulting in a resource state or subspace which is always available.

Arguably, the first examples of dissipation engineering in atomic physics were optical pumping and laser cooling. The former is a dissipation technique to initialise a qubit register in a separable state. The latter is applied to a boson (quantum harmonic oscillator) and creates a thermal state of motion. However, the term “dissipation engineering” is more commonly used to describe techniques that rely on quantum interference to generate non-classical states. In the oscillator state-space, such methods have been developed only more recently, and used create and stabilise novel quantum states of motion, such as squeezed vacuum states [Kie15], Schrödinger cat states [Lo15; Ger21] and grid states [dNee20].

In the qubit realm, the last decade has seen a range of proposals and experiments demonstrating dissipative entanglement generation. The methods employed for this purpose can be broadly categorised into three classes<sup>2</sup>:

1. **Resonance engineering.** With a combination of precisely targeted Stark-shift drives and global repumping drives, one can engineer a novel energy-level diagram. An additional probe field can then induce a resonant process whereby the system is pumped into the desired state, while the leakage processes out of the desired state are off-resonant. Such methods were theoretically proposed for entangling two atoms in optical cavities [Kas11; Rei12] and later used experimentally to entangle two trapped ions [Lin13] and two superconducting qubits [Sha13; Kim16]. In the trapped-ion implementation, an additional “coolant” ion was necessary for the protocol, as the repumping drive involved sympathetic cooling. Resonance engineering was also used in theoretical proposals for multi-qubit entangled state generation [Cho11; Lin16; Rei16; Col21] and quantum error correction [Coh14; Rei17]. The downside of resonance engineering is its speed. High-fidelity operation requires the off-resonant processes to be strongly attenuated, which in turn necessitates that the pumping rate is significantly lower than the resonance splittings. This results in long preparation times for the desired states [Hor18; Dou20; Col21].
2. **Symmetry engineering.** This method is similar to resonance engineering, with the steady-state obtained as a dark-state of a coupling Hamiltonian and repumping drives. However, leakage out of the target state is suppressed not due to a frequency selectivity, but due to destructive interference. Such a scheme was theoretically proposed for entangling two trapped ions without a need for a “coolant” ion in [Ben14], but to my knowledge never experimentally implemented. A similar scheme was proposed in [Hor18] to entangle two hyperfine  $^9\text{Be}^+$  qubits, and is being implemented as NIST. Both methods exhibit shorter convergence timescales compared to resonance

---

<sup>2</sup> Thanks to F. Reiter for introducing me to those helpful categories

engineering methods, and thus should result in higher steady-state fidelities. However, both schemes suffer from susceptibility to ion heating and temperature.

3. **Dissipative circuit engineering.** Here, the desired target state is created by optimising a quantum circuit that includes both unitary and dissipative elements. This was used in a trapped-ion system in [Bar11] to create two- and four-qubit entanglement. In both cases, an ancilla qubit was necessary to mediate the dissipative step.

The distinction between the different classes is not always sharp. Resonance engineering techniques do rely on certain symmetries to separate the resonances, and the target state is usually an antisymmetric singlet state. Furthermore, many of the schemes can be implemented either as a continuous dissipative process, or by interleaving periods of dissipative and unitary evolution. Thus, circuit engineering could be considered a method of resonance engineering.

In this chapter, we demonstrate a novel method of two-qubit entangled state generation based on symmetry engineering. The technique uses a “collective optical pumping” process which does not couple to the target Bell state  $|\Psi^-\rangle \equiv (|\uparrow\downarrow\rangle - |\downarrow\uparrow\rangle)/\sqrt{2}$ . At a high level, the protocol can be applied in any platform where suitable couplings are available, and requires only global fields which equally address each physical system. In the trapped-ion implementation, the collective drive Hamiltonian can be based on the MS interaction. Thanks to its first-order temperature insensitivity, the protocol does not require ground-state cooling or sympathetic cooling. Furthermore, it can be readily implemented in any trapped-ion setup where both optical and Zeeman qubit drives are available.


The chapter is structured as follows. In Sec. 9.1 we present our protocol in detail, including a high-level platform-independent description, as well as the specific implementation in  $^{40}\text{Ca}^+$ . We then present a theoretical model of the sequence. In Sec. 9.2 we simulate the protocol’s convergence rate and error sensitivity. Sec. 9.3 describes the experimental challenges involved in implementing the protocol, and how the ideal protocol was adapted in the face of unanticipated error channels. Finally, Sec. 9.4 presents the experimental measurements of the protocol’s performance and its error-susceptibility.

The main results presented in this chapter are now being prepared for publication as [Mal21].

## 9.1. Protocol overview

In this section, we present the theoretical description of our protocol. Sec. 9.1.1 presents the main idea behind the scheme in a platform-agnostic fashion. In Sec. 9.1.2, we present the methods for modelling the protocol mathematically in either a continuous or discrete fashion. Finally, Sec. 9.1.3 presents the implementation on a pair of trapped  $^{40}\text{Ca}^+$  ions.

## 9.1.1. High-level description



**Figure 9.1:** High-level description of the protocol. Coupling (A) is a collective excitation drive from  $|\downarrow\downarrow\rangle$  to  $|ee\rangle$ . Drive (B) is a repump channel that repopulates both qubit ground states. Drive (C) is a global single-qubit rotation with shuffles the population between the singlet states in the ground-state manifold. When drives (A)-(C) are switched on, the system is pumped into a maximally entangled state  $|\Psi^-\rangle$ .

The scheme uses a combination of coherent and incoherent drives applied to two coupled multi-level systems to create a maximally entangled state. The scheme we use is illustrated using the relevant states of two multi-level systems in Fig 9.1 a). We consider a ground “spin” state manifold consisting of the collective spin states  $|\downarrow\downarrow\rangle$ ,  $|\uparrow\uparrow\rangle$ ,  $|\Psi^+\rangle \equiv (|\uparrow\downarrow\rangle + |\downarrow\uparrow\rangle)/\sqrt{2}$  (spin triplet) and  $|\Psi^-\rangle$  (spin singlet), as well as excited states, of which the most important for our purposes consists of both systems in a particular excited state  $|e\rangle$ . Three primary elements define the pumping process. The first is a collective excitation (A) from the state  $|\downarrow\downarrow\rangle$  to the doubly excited state  $|ee\rangle$ . Its collective nature means that it does not couple to the other states in the ground-state manifold. The state  $|ee\rangle$  is quenched through a decay channel (B) that acts independently on the qubits, thus redistributing population from  $|ee\rangle$  into all spin states. The combination of these two provides a collective pumping that moves the ground state population from  $|\downarrow\downarrow\rangle$  to the other ground states. To prepare only the singlet, an additional symmetric drive (C) is used which rotates both qubits using a common angle and phase. Due to its symmetry, this drive cycles population within the triplet subspace, while leaving the singlet untouched. Thus the triplet states have a chance of being emptied through the collective pumping, while the population in the singlet is dark to all drives.  $|\Psi^-\rangle$  then becomes the steady-state of the system.



### 9.1.2. Theoretical model

The protocol can be implemented in a continuous manner or using a sequential application of each component. To model the former, we describe the drives in terms of their Hamiltonians and Lindblad jump operators. While the sequential method can also be described in the continuous formalism, its can be considerably simplified by modelling individual steps as unitary or dissipative maps.

#### Continuous model

In the continuous case, the drives (A)-(C) are switched on at the same time. We model the subsequent state evolution though a Lindblad master equation (Eq. (2.17)). The total interaction Hamiltonian  $H_I = H_A + H_C$  is the sum of the interactions Hamiltonians for drives (A) and (C) respectively. We model the collective excitation Hamiltonian (A) as:

$$H_A = \hbar J S_{x,e}^2, \quad (9.1)$$

where  $J$  is the effective collective excitation rate and

$$S_{x,e} = \sigma_{x,\downarrow e} \otimes \mathbf{1} + \mathbf{1} \otimes \sigma_{x,\downarrow e}.$$

where  $\sigma_{x,\downarrow e} = |e\rangle\langle\downarrow| + |\downarrow\rangle\langle e|$  is the Pauli-X operator acting on the  $\{|\downarrow\rangle, |e\rangle\}$  subspace, and  $\mathbf{1}$  is a  $3 \times 3$  identity operator. Note that for a two-level system, i.e. if  $|\uparrow\rangle$  was not part of the Hilbert space,  $\sigma_x^2 = \mathbf{1}$  and hence  $H_A$  reduces to an effective  $\sigma_x \otimes \sigma_x$  coupling. However, in the qutrit space  $\sigma_{x,\downarrow e}^2 \neq \mathbf{1}$  and Eq. (9.1) is not reducible to a pure  $\sigma_{x,\downarrow e} \otimes \sigma_{x,\downarrow e}$  coupling.

The Hamiltonian of the ground-state drive (C) is the Rabi oscillation Hamiltonian from Eq. (2.18) applied to both qubits at once. Taking for simplicity<sup>3</sup>  $\delta = 0$  and  $\phi = 0$ , we obtain

$$H_C = \hbar \frac{\Omega_C}{2} S_x$$

where

$$S_x = \sigma_x \otimes \mathbf{1} + \mathbf{1} \otimes \sigma_x$$

where  $\sigma_x = |\uparrow\rangle\langle\downarrow| + |\downarrow\rangle\langle\uparrow|$  is the ground state qubit Pauli-X operator. Note that we have assumed both qubits are driven with equal Rabi frequency. This symmetry is necessary for  $|\Psi^-\rangle$  to be unaffected by  $H_C$ .

---

<sup>3</sup> None of these assumptions is necessary. Non-zero detuning  $\delta$  reduces the effective population transfer efficiency, in turn slowing down the protocol. On the other hand, the choice of  $\phi$  does not influence the convergence rate.

## 9. OPTICAL PUMPING INTO A MAXIMALLY ENTANGLED STATE

The dissipative drive (B) is modelled by four jump operators:

$$\begin{aligned} L_{e \rightarrow \downarrow}^{(1)} &= \sqrt{p_{e \rightarrow \downarrow} \kappa} \sigma_{-, \downarrow e} \otimes \mathbf{1} & L_{e \rightarrow \uparrow}^{(1)} &= \sqrt{p_{e \rightarrow \uparrow} \kappa} \sigma_{-, \uparrow e} \otimes \mathbf{1}, \\ L_{e \rightarrow \downarrow}^{(2)} &= \sqrt{p_{e \rightarrow \downarrow} \kappa} \mathbf{1} \otimes \sigma_{-, \downarrow e} & L_{e \rightarrow \uparrow}^{(2)} &= \sqrt{p_{e \rightarrow \uparrow} \kappa} \mathbf{1} \otimes \sigma_{-, \uparrow e}. \end{aligned}$$

where  $\sigma_{-, \downarrow e} = |\downarrow\rangle\langle e|$  and  $\sigma_{-, \uparrow e} = |\uparrow\rangle\langle e|$ ,  $\kappa$  is the decay rate and  $p_{e \rightarrow \downarrow}$  ( $p_{e \rightarrow \uparrow}$ ) represent the probability that an ion decays from  $|e\rangle$  to  $|\downarrow\rangle$  ( $|\uparrow\rangle$ ). We assume that no decay occurs outside of the three-level subspace, and thus  $p_{e \rightarrow \downarrow} + p_{e \rightarrow \uparrow} = 1$ . This allows us to parametrise the branching ratio through:

$$\begin{aligned} p_{e \rightarrow \downarrow} &= \sin^2(\gamma), \\ p_{e \rightarrow \uparrow} &= \cos^2(\gamma). \end{aligned} \tag{9.2}$$

Note also that, unlike simultaneously applied Hamiltonians, individual jump operators of simultaneously occurring decay processes cannot be added together to form a single jump operator. The evolution of the system can be simulated by solving the Lindblad master equation with  $H_I = H_A + H_C$  and the jump operators  $L_j^{(i)}$ , and is parametrised by  $(J, \Omega_C, \kappa, \gamma)$ . Note that, regardless of the parameter choice, the singlet state  $|\Psi^-\rangle$  is the dark state of every individual Hamiltonian and every jump operator:

$$\begin{aligned} H_A |\Psi^-\rangle &= H_C |\Psi^-\rangle = 0 \\ L_{e \rightarrow \downarrow}^{(i)} |\Psi^-\rangle &= L_{e \rightarrow \uparrow}^{(i)} |\Psi^-\rangle = 0. \end{aligned}$$

Thus,  $|\Psi^-\rangle$  is indeed a dark state when drives (A), (B) and (C) are switched on simultaneously.

### Sequential protocol

When drives (A)-(C) are applied sequentially, the analysis can be simplified, avoiding the need to solve dynamical equations. Consider a single cycle of the protocol, where the drive (A) is first applied for duration  $t_A$ , followed by the drive (B) of duration  $t_B$  and the drive (C) of duration  $t_C$ . The first step implements a unitary:

$$U_A(\Phi) = e^{-i\Phi S_{x,e}^2}, \tag{9.3}$$

where  $\Phi = Jt_A$ , obtained by integrating  $H_A$  in Eq. (9.1). This unitary provides a full population transfer from  $|\downarrow\downarrow\rangle$  to  $|ee\rangle$  for  $\Phi = \pi/4$ .

For the second step, we assume for simplicity that  $t_B \gg 1/\kappa$ , and thus the population is fully repumped into the ground-state manifold. As a result, it is no longer necessary to keep track of  $\kappa$ , and the repump process can be fully parametrised by  $\gamma$ . The repump step

is then modelled by a quantum operation  $\varepsilon_B$ , defined as

$$\varepsilon_B(\rho) = \sum_{i=1}^3 \sum_{j=1}^3 E_{i,j} \rho E_{i,j}^\dagger$$

with the nine Kraus maps given by

$$E_{i,j} = E_i \otimes E_j$$

where

$$\begin{aligned} E_0 &= |\downarrow\rangle\langle\downarrow| + |\uparrow\rangle\langle\uparrow| \\ E_1 &= \sqrt{p_{e \rightarrow \downarrow}} \sigma_{-, \downarrow e} \\ E_2 &= \sqrt{p_{e \rightarrow \uparrow}} \sigma_{-, \uparrow e}, \end{aligned}$$

with the branching ratio parametrised as in Eq. (9.2).

The third step of every cycle corresponds to a unitary

$$U_C(\theta) = e^{i\frac{\theta}{2}\sigma_x} \otimes e^{i\frac{\theta}{2}\sigma_x}$$

obtained by integrating  $H_C$  in Eq. 9.1.2 with  $\theta = \Omega_C t_C$ .

While a sequence of unitary operations can be modelled by multiplying their individual unitary matrices, this is not the case for Kraus maps. In order to simplify the analysis, we turn to the superoperator formalism [Woo11]. In the superoperator formalism, states in  $d$ -dimensional Hilbert space are represented as vectors of length  $d^2$ . Formally, a state vector  $\vec{\rho}$  is obtained by column-stacking the density matrix, i.e. for a density matrix given by

$$\rho = \sum_{i=1}^d \sum_{j=1}^d \rho_{i,j} |i\rangle\langle j|$$

the vector representation is

$$\vec{\rho} = (\rho_{1,1}, \rho_{1,2}, \dots, \rho_{1,d}, \rho_{2,1}, \rho_{2,2}, \dots, \rho_{d,d})^T.$$

Any quantum operation  $\varepsilon$  can be written in the Kraus representation as

$$\varepsilon(\rho) = \sum_k E_k \rho E_k$$

Then, the superoperator representation of  $\varepsilon$  is given by:

$$S_\varepsilon = \sum_k E_k^* \otimes E_k, \tag{9.4}$$

## 9. OPTICAL PUMPING INTO A MAXIMALLY ENTANGLED STATE

where  $E^*$  is a matrix obtained by taking a complex conjugate of every element of  $E$ . Now, the effect of the operation is to transform an initial state  $\vec{\rho}$  into a state  $\vec{\rho}'$  given by

$$\vec{\rho}' = S_\varepsilon \vec{\rho},$$

i.e. the transformation can be modelled as a single matrix multiplication, albeit of significantly larger matrices [Gre15]. As a result, the superoperator of a sequence of transformations is the matrix product of the superoperators of the individual steps. Note that Eq. (9.4) implies that the superoperator of a unitary transformation  $U$  is simply  $S_U = U^* \otimes U$ . In the superoperator formalism, we can calculate the state of the system after  $N$  cycles of the protocol as

$$\vec{\rho}_N = \left[ S(\Phi, \gamma, \theta) \right]^N \vec{\rho}_0, \quad (9.5)$$

where  $\vec{\rho}_N$  and  $\vec{\rho}_0$  are the vector representation of the final and initial density matrices respectively, and the single-cycle superoperator is given by

$$S(\Phi, \gamma, \theta) = S_C(\theta) S_B(\gamma) S_A(\Phi), \quad (9.6)$$


where  $S_A(\Phi)$ ,  $S_B(\gamma)$  and  $S_C(\theta)$  are the superoperators for the drives (A), (B), and (C) respectively. In analogy to the continuous implementation, the vector representation of the singlet state  $|\Psi^-\rangle$  is a +1 eigenstate of  $S(\Phi, \gamma, \theta)$  regardless of the parameter settings, and thus is a dark state of the sequential protocol.

### 9.1.3. Trapped-ion implementation in $^{40}\text{Ca}^+$

The protocol can be implemented on a pair of  $^{40}\text{Ca}^+$  ions. To that end, we take advantage of the hybrid encoding introduced in Sec. 2.2.7. The qubit states are encoded into the Zeeman sub-levels of the electronic ground state, denoted as  $|\downarrow\rangle = |S_{1/2}, m_j = -1/2\rangle$  and  $|\uparrow\rangle = |S_{1/2}, m_j = +1/2\rangle$ . The ancillary state is selected among the  $D_{5/2}$  Zeeman sub-levels. In our experiments, we chose  $|e\rangle = |D_{5/2}, m_j = -1/2\rangle$ , because the corresponding optical qubit transition  $|\downarrow\rangle \leftrightarrow |e\rangle$  has the lowest sensitivity to B-field fluctuations. The drives used to implement the steps (A), (B) and (C) are shown in Fig. 9.2.

The collective excitation step (A) was implemented using an MS interaction on the optical qubit transition (Fig. 9.2, red). Consider a bichromatic MS drive incident on two ions. Adapting the results in Sec. 2.4 to a three-level system, and choosing  $\phi_s = \phi_m = 0$  for simplicity, we write the interaction Hamiltonian as

$$H_{\text{MS}} = \frac{\eta \hbar \Omega}{2} S_{x,e} (\hat{a} e^{i\delta t} + \hat{a}^\dagger e^{-i\delta t})$$



**Figure 9.2:** Atomic transitions and drives in  $^{40}\text{Ca}^+$  used in the implementation described in this work. Drive (A) is implemented using a bichromatic  $\lambda = 729\text{ nm}$  laser which implements an MS hamiltonian with sideband detuning  $\delta$ . Drive (B) uses laser light at  $\lambda = 854\text{ nm}$  and  $\lambda = 866\text{ nm}$  to couple  $|e\rangle$  and  $D_{3/2}$  to a short-lived excited state  $P_{3/2}$ , causing a decay into both ground states in the  $S_{1/2}$  manifold. Drive (C) is a resonant global Zeeman qubit rotation implemented by passing current through a track on the trap carrier PCB.

and the corresponding unitary as

$$U_{\text{MS}}(t) = D(\alpha(t)S_{x,e})e^{-i\Phi(t)S_{x,e}^2} \quad (9.7)$$

where  $\alpha(t)$  is an oscillator phase-space displacement amplitude (Eq. (2.28)) and  $\Phi(t)$  is a collective phase factor (called  $\beta(t)$  in Eq. (2.30)).

The MS coupling can be reduced to an effective  $S_{x,e}^2$  Hamiltonian in two limits. The first, appropriate to either a continuous or step-wise implementation, is when

$$\delta \gg \eta\Omega \quad (9.8)$$

and so the oscillator excitation can be neglected [Kim10]. Then  $|\alpha(t)| \ll 1$  and the unitary can be approximated as

$$U_{\text{MS}}(t) \approx e^{i\Phi(t)S_{x,e}^2}$$

which corresponds to an effective Hamiltonian of

$$H_{\text{eff}} = \hbar J S_{x,e}^2 \quad (9.9)$$

## 9. OPTICAL PUMPING INTO A MAXIMALLY ENTANGLED STATE

where

$$J = \frac{\partial \Phi}{\partial t} \approx \frac{\eta^2 \Omega^2}{4\delta}. \quad (9.10)$$

Indeed, the effective Hamiltonian  $H_{\text{eff}}$  in Eq. (9.9) can be used to implement the Hamiltonian  $H_A$  in Eq. (9.1), which is a building block of the continuous implementation of the protocol.

The second limit, appropriate only to a sequential implementation, is to apply the MS coupling for a time  $t$  chosen such that  $\alpha(t) = 0$ . As discussed in Sec. 2.4, this is achieved when  $t = 2m\pi/\delta$ , leading to

$$U_{\text{MS}}(t) = e^{i\Phi(t)S_{x,e}^2} \quad (9.11)$$

where  $\Phi(t) = \eta^2 \Omega^2 m\pi / (2\delta^2)$  and  $m \in \mathbf{Z}$ . Thus, the MS unitary in Eq. (9.11) can be used to implement the unitary for the drive (A) in Eq. (9.3), which is a building block of the sequential implementation of the protocol.

While the single-loop MS gate corresponds to  $m = 1$  and  $\Phi(t) = \pi/8$ , full population transfer from  $|\downarrow\downarrow\rangle$  to  $|ee\rangle$  requires  $\Phi = \pi/4$ , which is achieved for  $m = 2$ . Thus, the drive (A) physically corresponds to an MS gate applied twice in a row. As discussed in Sec. 2.4.3, the performance of any multi-loop sequence can be improved by phase modulation. Thus, in the simulations and experiments discussed below, I always assume that  $U_{\text{MS}}(t)$  is implemented as two segments of equal motional phases  $\phi_{m1} = \phi_{m2}$ , but with a spin phase offset of  $\phi_{s2} = \phi_{s1} + \pi$ . Such phase modulation reduces errors caused by spin-motion coupling both in simulation and experiment (Sec. 9.3.1). Note that the phases  $\phi_{s1}$  and  $\phi_{m2}$  for the first pulse in any cycle are usually set to zero in the simulation, but can be arbitrary.

Both limits correspond to the same effective Hamiltonian  $H_{\text{eff}}$  with the coupling rate  $J$  given by Eq. (9.10). However, the continuous limit in Eq. (9.8) implies that  $J \ll \eta\Omega/4$ . At the same time, no such restriction applies to the pulsed implementation, which is usually performed with  $J = \eta\Omega/8$ . Therefore, we expect the sequential protocol to be faster and more robust to errors. As a result, the remainder of the theoretical discussion in this chapter, as well as the experimental implementation, are focused on the step-wise protocol only. The continuous protocol was analysed by I. Rojko, and a detailed discussion of its performance, convergence rate and error-resilience can be found in [Mal21, Supp. Mat.]

The remaining drives are implemented identically regardless of the implementation. Repump (B) is implemented by turning on the dipole lasers at  $\lambda = 854 \text{ nm}$  and  $\lambda = 866 \text{ nm}$  (Fig. 9.2, blue). The former drives the primary repump channel by coupling all the  $D_{5/2}$  sub-levels to the  $P_{3/2}$  state, which subsequently decays primarily into the ground state manifold (Sec. 2.2.1). The latter laser repumps the population in  $D_{3/2}$ , which would otherwise become gradually populated<sup>4</sup> due to a small branching ratio from  $P_{3/2}$  to  $D_{3/2}$  (Tab. 2.1). The symmetric drive (C) is implemented as a Zeeman single-qubit Rabi drive

<sup>4</sup> In the first attempt at the implementation of the collective optical pumping protocol by V. Negnevitsky [Neg18b], this drive was indeed forgotten, which was one of the reasons for large errors

(Sec. 5.5.3 and Fig. 9.2, green), created by passing oscillating current through a track on the trap carrier PCB (Sec. 3.4.1).

## 9.2. Performance analysis

In this section, we analyse the performance of an idealised discrete protocol. We begin by exploring the general properties of the superoperator  $S(\Phi, \gamma, \theta)$  in Eq. (9.6). Formally, state  $\vec{\rho}$  is a steady-state of the protocol if and only if

$$S(\Phi, \gamma, \theta)\vec{\rho} = \vec{\rho},$$

i.e. when  $\vec{\rho}$  is a +1 eigenvector of  $S(\Phi, \gamma, \theta)$ . It is easy to verify that, for any values of  $(\Phi, \gamma, \theta)$ , the vector corresponding to the singlet state  $|\Psi^-\rangle$  is indeed a +1 eigenvector of  $S(\Phi, \gamma, \theta)$ , and thus a steady state. However, in order to assert convergence towards  $|\Psi^-\rangle$ , we require that it is also the *only* steady state. Through numerical verification, we determine that this is indeed the case for “almost all” parameter settings, i.e. for all values  $(\Phi, \gamma, \theta)$  except a finite set of exceptions:

1. The collective excitation (A) exhibits an exception at  $\Phi = m\pi/2$ , where  $m \in \mathbf{Z}$ , i.e. no population transfer. In this case, there are four steady states in the ground-state manifold.
2. The decay step (B) exhibits an exception at  $\gamma = \frac{m\pi}{2}$ , where  $m \in \mathbf{Z}$ . For even  $m$ , corresponding to  $p_{e \rightarrow \downarrow} = 0$ , the ground state  $|\uparrow\uparrow\rangle$  is also a steady state, while for odd  $m$ , which corresponds to  $p_{e \rightarrow \downarrow} = 0$ , the ground state  $|\downarrow\downarrow\rangle$  is a steady state
3. The symmetric drive (C) exhibits an exception at  $\theta = m\pi$  where  $m \in \mathbf{Z}$ . For even  $m$ ,  $|\uparrow\uparrow\rangle$  is a steady state, while for odd  $m$ , the steady state is  $|\Psi^+\rangle$ .

The impact of these exceptions is a slow-down of protocol convergence near the exception points. However, their experimental impact is negligible, and the singlet state  $|\Psi^-\rangle$  is, for all intents and purposes, the only true steady state of our protocol.

### 9.2.1. Convergence rate

The effect of different choices of  $(\Phi, \gamma, \theta)$  on the convergence rate can be studied in two ways. Numerically, we can simulate the protocol with Eq. (9.5) for different parameter settings and different initial states. We define the singlet error (or infidelity) after  $N$  cycles as

$$\epsilon = 1 - F(|\Psi^-\rangle)$$

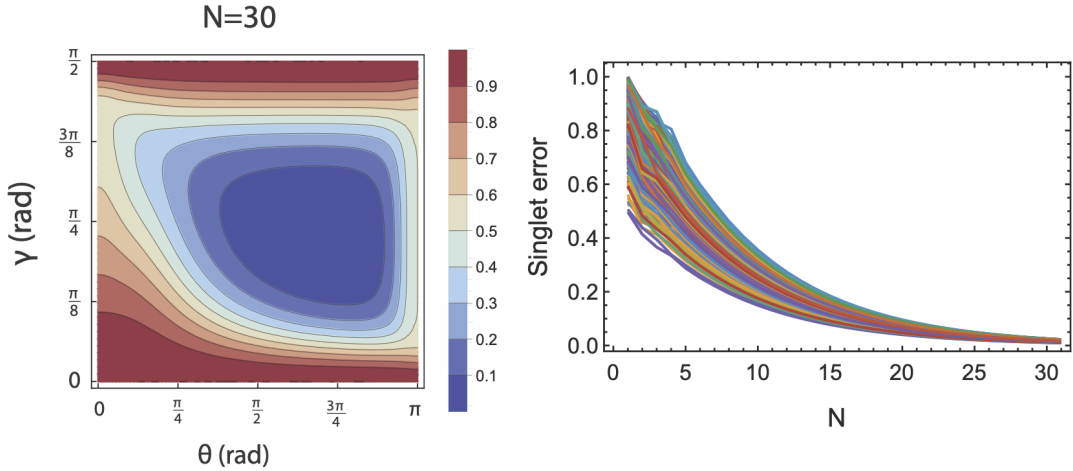
where

$$F(|\Psi^-\rangle) = |\langle \Psi^- | \rho | \Psi^- \rangle|.$$

## 9. OPTICAL PUMPING INTO A MAXIMALLY ENTANGLED STATE

While  $\epsilon \rightarrow 0$  as  $N \rightarrow \infty$ , the optimal settings are those that guarantee the fastest convergence, i.e. minimise  $\epsilon$  for a fixed  $N$ .

It is easy to verify that  $\Phi = \pi/4$  (i.e. full state transfer from drive A) always produces the fastest convergence. We find the optimal values of  $\gamma$  and  $\theta$  by simulating  $\epsilon$  for  $N = 30$  cycles and an initial state  $|\downarrow\downarrow\rangle$ . This is shown in Fig. 9.3 (left). We find that there is a relatively broad region where convergence occurs at a comparable rate, indicating that  $\gamma$  and  $\theta$  do not require precise calibration. With a numerical search, we find  $\epsilon$  to be the lowest for  $\theta \approx 0.72\pi$  and  $\gamma \approx 0.22\pi$ .



**Figure 9.3:** Simulations of the protocol convergence rate. (Left) Singlet error  $\epsilon$  after  $N = 30$  cycles of the protocol acting on an initial state  $|\downarrow\downarrow\rangle$  and the drive (A) set to the optimal value  $\Phi = \pi/4$ . The lowest error corresponds to  $\theta = 0.72\pi$  and  $\gamma = 0.22\pi$  (Right) Simulated singlet error  $\epsilon$  as a function of the number of cycles  $N$  for range separable initial states generated at random. For the optimal parameters as simulated here, each curve asymptotically converges to an exponential decay trace  $\epsilon = 1 - C_0 e^{-N/N_0}$  with  $N_0 = 7.6$  cycles.

We study the convergence rate and its variability at the optimal settings by simulating the process with  $\theta = 0.72\pi$  and  $\gamma = 0.22\pi$  for a range of initial states. We generate 1000 random input density matrices corresponding to two-qubit separable pure states. For each input state, we simulate the ideal protocol and fit the fidelity with a function

$$\epsilon(N) = 1 - C_0 e^{-N/N_0}.$$

The simulated traces are shown in Fig. 9.3 (right). To find the average convergence rate, we ignore the data with  $N < 10$ , where  $\epsilon(N)$  is dominated by the initial state. This produces a convergence rate estimate of

$$N_0 = 7.6 \text{ cycles.}$$

The standard deviation of  $N_0$  for different initial states is at the level of  $1 \times 10^{-4}$ , indicating



that for  $N > 10$  the convergence rate is independent of the initial state. By comparison, fitting all the data results in  $N_0 = 7.7(2)$  cycles.

The convergence rate can be found theoretically through superoperator eigenvalue analysis. Recall that the steady-state of the process is given by the  $+1$  eigenvalue of  $S(\Phi, \gamma, \theta)$ . The remaining eigenstates of the superoperator all have eigenvalues with an absolute value less than 1, which describe their decay rates. The detailed analytical study of the problem was performed by I. Roikov, and is discussed in (cite our Arxiv paper). The summary of his findings is as follows. After the initial transients, the infidelity is given by

$$\epsilon(N) \approx C_0 \lambda_{\max}^N = C_0 e^{-N/N_0},$$

where  $C_0$  is the initial infidelity,  $N_0 = \log(\lambda_{\max})$  and  $\lambda_{\max}$  is the second largest eigenvalue of  $S$ . This eigenvalue is maximised, and hence the convergence is the fastest, for  $\theta \approx 0.72\pi$  and  $\gamma \approx 0.22\pi$ , exactly as we found by a numerical search. At those settings,  $\lambda_{\max}$  can be analytically calculated to equal  $\lambda_{\max} \approx 0.877$ , corresponding to  $N_0 \approx 7.6$  cycles. Thus, the analytical formalism fully explains the simulated performance.

### 9.2.2. Effect of errors

We now turn our attention to the question of the error-resilience of the protocol. In this section, we study this problem at a high level of abstraction, considering errors as quantum maps. While the choice of the maps is motivated by experimentally relevant noise sources, the description is general and platform-agnostic. The trapped-ion specific discussion is presented in the subsequent section 9.2.3, where we simulate the protocol dynamics including realistic error sources.

We initially focus on errors introduced by or during the drive (A), which is experimentally the most demanding. We thus model each cycle of the protocol with a superoperator

$$S(\Phi, \gamma, p, \theta) = S_C(\theta)S_B(\gamma)S_E(p)S_A(\Phi)$$

where  $S_E(p)$  is the superoperator representing an error, and  $p$  parametrises the error probability. For  $p > 0$ , we can find the steady-state singlet error  $\epsilon$  analytically by calculating the overlap of the  $+1$  eigenvector of  $S(\Phi, \gamma, p, \theta)$  with  $|\Psi^-\rangle$ , or numerically by repeatedly applying  $S(\Phi, \gamma, p, \theta)$  to any initial state. We now consider several specific error models.

#### Two-level depolarising errors

In the trapped-ion implementation, the drive (A) acts on the optical qubit only and is spectroscopically decoupled from  $|\uparrow\rangle$ . Under this assumption, we can describe the elementary errors introduced by (A) through elementary error channels acting only on the  $\{|\downarrow\rangle, |e\rangle\}^{\otimes 2}$

## 9. OPTICAL PUMPING INTO A MAXIMALLY ENTANGLED STATE

subspace. There are 16 elementary error channels, defined as

$$M_{i,j} = e^{-i(\pi/2)\sigma_{i,\downarrow e} \otimes \sigma_{j,\downarrow e}}$$

for  $i, j \in \{0, x, y, z\}$ , using  $\sigma_0 = \mathbf{1}$ . This operator definition extend the usual notion of single-qubit errors onto a three-level system. We can divide these Kraus maps into two classes:  $M_{i,j}$  is considered a *local* operator when  $i = 0$  or  $j = 0$ , and a *global* operator otherwise.

To study the influence of each of these channels individually, we simulate the effect of an error in a manner similar to a depolarising channel

$$\varepsilon_{i,j}(\rho) = (1 - p)\rho + pM_{i,j}\rho M_{i,j}^\dagger$$

acting once per cycle after drive (A). Later in the text, we refer to those errors with a shorthand notation, e.g.  $\varepsilon_{0,x}$  is an “IX” error and  $\varepsilon_{x,z}$  is an “XZ” error. For every such error, we estimate the steady-state error  $\epsilon(p)$  by fitting the numerical results for  $p < 0.01$ . The resulting errors are illustrated in Fig. 9.4, rounded to the nearest integer. We find that

	I	X	Y	Z
I	0	5p	5p	4p
X	5p	0	0	0
Y	5p	0	0	0
Z	4p	0	0	0

**Figure 9.4:** Visualisation of the steady-state error caused by a elementary depolarising errors acting with probability  $p$  per cycle on the two-level system formed by  $|\downarrow\rangle$  and  $|e\rangle$ . Different depolarising error channels are referred to through a short-hand notation of  $M_{i,j}$ . Steady-state error  $\epsilon$  for a depolarising map  $\varepsilon_{i,j}$  is shown in row  $i$  and column  $j$ .

the final fidelity is independent of all global errors. This is expected, as these maps only affect the system when both particles are in the  $\{|\downarrow\rangle, |e\rangle\}$  subspace, thus do not provide a leakage channel out of  $|\Psi^-\rangle$ :

$$\varepsilon_{i,j}(|\Psi^-\rangle\langle\Psi^-|) = |\Psi^-\rangle\langle\Psi^-| \quad \text{for global errors.}$$

On the other hand, local errors become amplified. Local bit-flip errors  $\varepsilon_{x,0}, \varepsilon_{y,0}, \varepsilon_{0,x}$  and

$\varepsilon_{y,0}$  create a steady-state error of

$$\epsilon \approx 5.2p \quad \text{local bit-flip error}$$

while local phase-flip errors  $\varepsilon_{z,0}$  and  $\varepsilon_{0,z}$  result in

$$\epsilon \approx 4.1p \quad \text{local phase flip error.}$$

In both cases, the result represents a linear approximation valid only for small values of  $p \ll 1$ . The fact that the steady-state error  $\epsilon$  exceeds the per-cycle error  $p$  is also a general feature of dissipative protocols. We encountered this for example in Sec. 5.2.2, where the optical pumping state preparation error exceeded the off-resonant excitation probability.

### Correlated local errors

Depolarising errors do not provide the most realistic description of experimentally relevant error channels. In this subsection, we consider three examples of error maps which, while remaining general, provide a quantum operator description of error sources which may be relevant to a trapped-ion implementation. They constitute a class that I refer to as *correlated local errors*.

The first example is a *correlated bit-flip error* caused by residual spin-motion entanglement at the end of the MS drive. Consider an MS drive applied to a pair of ground-state cooled ions. We can write the initial density matrix of the system as

$$\rho = \rho_s \otimes |0\rangle_m \langle 0|_m$$

where  $\rho_s$  is the internal-state density operator, and the subscript  $m$  denotes the motional degree of freedom. Following an application of Eq. (9.7), we can write the internal state of the two ions at the end of the drive (A) in the operator-sum representation as

$$\rho'_s = \sum_{n=0}^{\infty} E_n \rho E_n^\dagger$$

where

$$E_n = \langle n|_m U_{MS}(t) |0\rangle_m.$$

If the spin and motion are perfectly disentangled at the end of the drive (A) ( $\alpha(t) = 0$ ), we find that  $E_0 = U_{MS}(t)$ , and all the other Kraus operators are zero. However, residual spin-motion entanglement results in non-zero values of  $E_n$  for  $n > 0$ .

It is possible to simplify and abstract the analysis by a) assuming  $\Phi = \pi/4$  and  $|\alpha(t)| \ll 1$ , then b) expanding the all the operators  $E_i$  up to the leading order in  $|\alpha|$ , and c) keeping only the two dominant maps  $E_0$  and  $E_1$ . However, the resulting operators are not normalised

## 9. OPTICAL PUMPING INTO A MAXIMALLY ENTANGLED STATE

$(E_0^\dagger E_0 + E_1^\dagger E_1 \neq \mathbf{1})$ , which produces unphysical results after the map was applied repeatedly in the simulation. In order to obtain physically meaningful outputs with only two Kraus operators, I instead created a synthetic operator  $E'_0$  which is equal to  $E_0$  up to the leading order in  $|\alpha|$ , but produces a normalised process. This allowed me to write the whole operation in terms of only two maps:

$$\begin{aligned} E'_0(\alpha) &= \frac{1-b_2}{2} \mathbf{1}_2 + \frac{-1-b_2}{2} (\sigma_{x,\downarrow e} \otimes \sigma_{x,\downarrow e}) + \frac{1-b_1}{2} \mathbf{1}_1 + \mathbf{1}_0 \\ E_1(\alpha) &= -i\alpha S_{x,e}, \end{aligned} \quad (9.12)$$

where  $b_2 = \sqrt{1-4|\alpha|^2}$ ,  $b_1 = \sqrt{1-2|\alpha|^2}$ , and

$$\begin{aligned} \mathbf{1}_2 &= |\downarrow\downarrow\rangle\langle\downarrow\downarrow| + |\downarrow e\rangle\langle\downarrow e| + |e\downarrow\rangle\langle e\downarrow| + |ee\rangle\langle ee| \\ \mathbf{1}_1 &= |\downarrow\uparrow\rangle\langle\downarrow\uparrow| + |e\uparrow\rangle\langle e\uparrow| + |\uparrow\downarrow\rangle\langle\uparrow\downarrow| + |\uparrow e\rangle\langle\uparrow e| \\ \mathbf{1}_0 &= |\downarrow\downarrow\rangle\langle\downarrow\downarrow| \end{aligned}$$

are the operators which “select” whether two, one, or zero particles are in the optical qubit subspace. Intuitively,  $E_1$  describes a “correlated bit-flip error” caused by residual spin-motion entanglement at the end of the drive ( $A$ ), while  $E'_0$  is a sum of global maps which do not induce protocol error even for  $|\alpha| > 0$ .

Before investigating the effect of this error on the protocol steady-state, we turn our attention to a second possible source of correlated bit-flip errors, namely spectator mode excitation. We model the effect of undesired mode coupling as a unitary carrier drive of the optical qubit transition:

$$U_x(\beta) = e^{i(\beta/2)\sigma_{x,\downarrow e}} \otimes e^{i(\beta/2)\sigma_{x,\downarrow e}} \quad (9.13)$$

with a small angle  $\beta \ll 1$ .

An observant reader may now ask why we consider this error at all since, as discussed in Sec. 8.3.2, off-resonant carrier excitation during an MS gate is efficiently suppressed by pulse shaping. While that is true,  $U_x(\beta)$  is a proxy for two effects that cannot be fully captured by considering each ion as a three-level system. First, as was noted in Sec. 8.3.3, trap potential drifts can result in the MS pulse driving an off-resonant transition from  $|\downarrow\rangle$  to another state  $|e'\rangle$  in the  $D_{5/2}$  manifold. Secondly, as will be discussed in Sec. 9.3.2, spectral crowding is a larger problem in this protocol than it is for an MS gate, as spectator transitions out of  $|\uparrow\rangle$  must be avoided as well. In the experiments in Sec. 9.3.2, we will see how constant parameter calibration is required to avoid these spectator transitions. There are in total ten possible spectator transitions for the drive ( $A$ ), each exciting either  $|\downarrow\rangle$  or  $|\uparrow\rangle$  into one of five possible sub-levels of  $D_{5/2}$ . However, any one of these errors causes leakage out of  $|\Psi^-\rangle$  with the same probability. Therefore, I decided to model the effect of any (or multiple simultaneous) spectator mode excitations as  $U_x(\beta)$ , which is approximate but simplifies the analysis considerably.

In order to compare the effect of an error caused by residual spin-motion entanglement of magnitude  $\alpha$  to that of a spectator mode excitation with area  $\beta$ , we require a common parametrisation. Let us define  $p$  as the probability that an operation causes a single-qubit excitation when applied on  $|\downarrow\downarrow\rangle$ , i.e. that the state is measured in  $|\downarrow e\rangle$  or  $|e\downarrow\rangle$ <sup>5</sup>. Using Eq. (6.11) with Kraus maps  $E'_0$  and  $E_1$  as defined in Eq. (9.12), we find  $p = 2|\alpha|^2$  due to residual spin-motion entanglement. On the other hand, Eq. (9.13) gives  $p = \beta^2/2$  for an off-resonant excitation. Therefore, a reasonable description of the Kraus maps for both processes in terms of  $p$  is:

$$\begin{array}{ll} \{E_0(\sqrt{p/2}), E_1(\sqrt{p/2})\} & \text{spin-motion entanglement,} \\ U_x(\sqrt{2p}) & \text{spectator excitation.} \end{array}$$

Simulations reveal that, for any given  $p$ , the steady-state error  $\epsilon$  is almost the same regardless of which of the two processes is being simulated. This is not surprising, as the first-order effect of either error is a leakage out of  $|\Psi^-\rangle$  with the same probability of  $p/2$  per cycle. Consequently, we do not need to study them independently, or even experimentally resolve one error from the other. From now on, we model all correlated bit-flip errors as  $U_x(\sqrt{2p})$ , and henceforth refer to them collectively as  $IX + XI$ .

The third correlated local error we consider is the *correlated phase-flip error*, also referred to as a  $IZ + ZI$  error. Such an error can occur, for example, when laser power fluctuations cause a common-mode AC Stark shift of the optical qubit. This can be modelled as

$$U_z(\beta) = e^{i(\beta/2)\sigma_{z,\downarrow e}} \otimes e^{i(\beta/2)\sigma_{z,\downarrow e}}, \quad (9.14)$$


and once again we pick  $\beta = \sqrt{2p}$  as the error parametrisation. Other realistic error sources cause other types of correlated phase-flip errors. For example, laser frequency errors cause a correlated phase-shift where both ground state qubits acquire the same relative phase with respect to  $|e\rangle$ , while B-field fluctuations cause a relative shift of all three levels. However, as we shall see below, the protocol is insensitive to any type of a correlated phase-flip, since  $|\Psi^-\rangle$  lives in a decoherence-free subspace. We performed the calculations with  $U_z(\beta)$  as defined in Eq. (9.14), but any alternative form could be used as well.

Fig. 9.5 shows the simulated steady-state error  $\epsilon$  for correlated bit-flip errors (blue dashed line) and correlated phase-flip errors (yellow dashed line), as well as a comparison with the depolarising model presented previously (solid lines). For correlated bit-flip errors, we find that the steady-state infidelity is slightly reduced compared to the uncorrelated bit-flip errors in the depolarising model, now reading:

$$\epsilon \approx 3.2p \quad \text{correlated bit-flip error.}$$

---

<sup>5</sup> When parametrised in this way,  $p$  can be easily measured as  $|\downarrow\downarrow\rangle$  can be easily prepared. However, the main physical figure of merit is the probability to leakage out of  $|\Psi^-\rangle$ , which is equal to  $p/2$  for both processes.



**Figure 9.5:** Comparison of steady-state errors associated with local “depolarising errors” (solid lines) and local correlated errors (dashed lines). For bit-flip errors (blue), the correlated model results in a steady state error  $\epsilon \approx 3.2p$ , which is less than the error of  $\epsilon \approx 5.2p$  obtained for either,  $\epsilon_{x,0}$ ,  $\epsilon_{0,x}$ , or an average of the two. For phase-flip errors (yellow), the correlated model produces no protocol errors ( $\epsilon = 0$ ), while  $\epsilon_{z,0}$  or  $\epsilon_{0,z}$  error produces  $\epsilon \approx 4.1p$ .

We will see in Sec. 9.4.4 that this error model captures well the experimental observations. Note that we cannot claim that correlations in the noise lead to an increased protocol fidelity, as the “correlated” and “uncorrelated” error models are qualitatively different. Furthermore, the observed error reduction is largely an artefact of the chosen map parametrisation.

On the other hand, correlated phase-flips cause no steady-state error:

$$\epsilon = 0 \quad \text{correlated phase-flip error,}$$

since  $|\Psi^-\rangle$  is an eigenstate of all collective phase shifts. In this case, we can claim that correlations in phase shifts leads to error suppression.

### 9.2.3. Error simulation

Here, we consider example errors of the collective excitation MS Hamiltonian (Sec. 8.3.2) and simulate their effect on the steady-state error  $\epsilon$ . To include the possibility of errors, as well as the effect of phase modulation, we use an MS interaction Hamiltonian given by

$$\hat{H}_A = \frac{\eta\hbar\Omega}{2} S_{e,\phi} (\hat{a} e^{i(\delta+\epsilon_m+\epsilon_q)t} e^{i\phi_m} + \hat{a}^\dagger e^{-i(\delta+\epsilon_m-\epsilon_q)t} e^{-i\phi_m}), \quad (9.15)$$


where  $\epsilon_m$  is the motional frequency error,  $\epsilon_q$  is the qubit frequency error and the total spin operator is a sum of two single-ion operators  $S_{e,\phi} = S_{e,\phi}^{(1)} + S_{e,\phi}^{(2)}$ , where

$$S_{e,\phi} = |e\rangle\langle\downarrow| e^{i\phi_s} + |\downarrow\rangle\langle e| e^{-i\phi_s}.$$

Note that in writing down Eq 9.15, we took a rotating-wave approximation to remove components oscillating faster than  $\delta$ , and thus also assumed that other transitions from  $|\downarrow\rangle$  and  $|\uparrow\rangle$  are spectroscopically decoupled. Furthermore,  $\hat{H}_A$  contain no dipole couplings, thus it neglects the effect of the dipole AC Stark shifts.

We simulate each cycle as a sequence of two rectangular pulses of length  $t$ . We set  $\phi_{s1} = \phi_{m1} = 0$  for the first loop while the second loop has  $\phi_{s2} = 0$ ,  $\phi_{m2} = \pi - \delta t$ , which implements “Aussie-style” phase modulation discussed in Sec. 2.4.3 [Mil18b]. The simulation parameters are:  $\delta = 2\pi \times 15$  kHz,  $t = 2\pi/\delta = 66.6$   $\mu$ s,  $\eta = 0.028$ ,  $\Omega = \delta/(2\eta) = 2\pi \times 265$  kHz. For comparison, we also simulate a single-loop MS gate (total length  $t$ ) and a two-loop MS gate (total length  $t \times \sqrt{2}$ ). For all the simulations, the state is initialised in  $|\downarrow\downarrow\rangle$  and in the motional ground state. For our protocol, we extract the fidelity with  $|\Psi^-\rangle$  after 80 cycles, while for the MS gates, we calculate the fidelity between the final state and  $(|\downarrow\downarrow\rangle - i|ee\rangle)/\sqrt{2}$ .

We simulate the effect of slow and shot-to-shot errors in the optical qubit frequency or the motional frequency by setting  $\epsilon_q > 0$  or  $\epsilon_m > 0$  respectively in Eq. (9.15). Slow and shot-to-shot fractional Rabi frequency error  $\epsilon_\Omega$  is simulated by setting  $\Omega = (1 + \epsilon_\Omega)\delta/(2\eta)$  in Eq. (9.15). Finally, we consider collective Markovian optical qubit dephasing by solving the Lindblad master equation with a jump operator  $L = \sqrt{\Gamma}(\sigma_{z,\downarrow e}^{(1)} + \sigma_{z,\downarrow e}^{(2)})$ . We assume the Markovian error is only acting during the drive (A).



**Figure 9.6:** Comparison of the Bell-state error of an MS gate (green, solid), phase-modulated two-loop MS gate (green, dashed) and 80 cycles of our protocol (blue) applied to an initial state  $|\downarrow\downarrow\rangle$ . Note that the ideal Bell state is a spin triplet for the MS gates, and a spin singlet for our protocol.

## 9. OPTICAL PUMPING INTO A MAXIMALLY ENTANGLED STATE

Fig. 9.6 shows the simulations of the Bell-state fidelities in the presence of these experimental errors. For the errors in the left column, we find our protocol to be beneficial over an MS gate. The reduced sensitivity to optical qubit frequency offsets can be understood as follows. Recall that the main effect of a qubit frequency error on an MS gate is that it affects the phase of the final entangled state (Fig. 8.17 in Sec. 8.3.2). Our protocol is insensitive to correlated phase errors and therefore achieves error suppression. Experimentally, the optical qubit frequency errors can be caused by drifts in magnetic fields or the qubit laser frequency.

The reduced sensitivity to the Rabi frequency errors is also expected since an error in  $\Omega$  only changes the collective excitation phase  $\Phi$  from the ideal value of  $\pi/4$ . Errors in  $\Phi$  do not cause a steady-state error  $\epsilon$ , but only slow down the convergence of the protocol. One way to look at an error in  $\Omega$  is as causing an unwanted  $\varepsilon_{x,x}$  depolarising process, which leaves the singlet fidelity unaffected. Indeed, the simulated error increase for 80 cycles is solely a result of a slow-down in convergence.

On the other hand, the protocol is not beneficial over an MS gate for the errors in the right column of Fig. 9.6. For gate mode frequency offsets, such as caused by drifting trap potentials or cross-Kerr coupling, we find  $\epsilon$  to be comparable with the MS gate error. This is because  $\epsilon_m > 0$  causes two effects of comparable magnitude on a single MS gate. First, the gate area is modified, corresponding to an undesired  $\epsilon_{x,x}$  depolarising error. Second, the phase-space loop does not close, resulting in residual spin-motion entanglement and hence a correlated bit-flip error. Our protocol removes the first error but amplifies the second, and thus the overall error stays comparable. While the simulations presented in Fig. 9.6 indicate an error suppression for small frequency errors, we note that this could be achieved equally well (but in a much shorter time) by using higher-order phase modulation.

Likewise, we find no benefit of using our protocol if an experiment is limited by fast (Markovian) optical qubit dephasing, which could be caused by fast laser or magnetic field noise. This is because dephasing during the MS gate causes a bit-flip error with probability  $p$  proportional to the gate time  $t$  (Eq. (8.13)). This error affects the dissipative protocol the most, since each loop contains an MS drive of length  $2 \times t$ , twice as long as a single-loop MS gate.

An important effect not included in the simulation is spontaneous emission from the state  $|e\rangle$ . The optical qubit lifetime  $T_1$  limits the Bell state error of an MS gate of length  $t$  to  $\approx 0.5t/T_1$ , where  $T_1 = 1.1$  s for  $^{40}\text{Ca}^+$  (Sec. 8.3.2). The dissipative protocol amplifies this error, resulting in steady-state infidelity of  $\epsilon \approx 1.5t_A/T_1$ , where  $t_A$  is the length of the drive (A).

Finally, the entangling drive (A) may not be the only source of errors in the protocol. The dissipative step (B) needs to be sufficiently spectroscopically decoupled to avoid exciting the ground state qubit. Besides this, it is only necessary to ensure the population does not get trapped outside of the protocol subspace. Ultimately, every protocol cycle event can cause heating due to photon recoil, but the average number of photon recoil events will be



significantly lower than the number of cycles. For a branching ratio parametrized by  $\gamma$ , the average number of photons scattered during the full protocol can be calculated to equal

$$\langle n_\gamma \rangle = 4 \csc(2\gamma)^2.$$

For  $\gamma = 0.22\pi$  this gives  $n_\gamma \approx 4$  photons. Thus, photon scattering during the protocol produces negligible heating, comparable to that caused by optical pumping.

The most important assumption about the drive (C) is that it couples to both ground-state qubits with equal Rabi frequencies. For a small imbalance of  $\Omega_C^{(2)} = \Omega_C^{(1)}(1 + \epsilon)$ ,  $\epsilon \ll 1$  and  $\theta \approx 0.72\pi$ , the singlet state gets depopulated with probability  $p \approx 1.4\epsilon^2$  per cycle. This increases the steady-state error similarly to a correlated bit-flip during the drive (A). Obtaining  $p \ll 10^{-4}$  requires  $\epsilon \ll 10^{-2}$ , which can easily be the case for long-wavelength radiation. Magnetic field noise does not impact the steady-state fidelity, since  $|\Psi^-\rangle$  lives in a decoherence-free subspace, but noise in the magnetic field gradient does cause coupling between  $|\Psi^-\rangle$  and  $|\Psi^+\rangle$ . AC Stark shifts caused by off-resonant couplings of drives (A)-(C) to spectator modes generally do not cause an error, since they affect  $|\downarrow\rangle$  and  $|\uparrow\rangle$  equally. Furthermore, any drives that affect the Zeeman qubit states differentially (from the  $\lambda = 729\text{ nm}$  laser) are still expected to couple to both ions equally, and thus leave  $|\Psi^-\rangle$  unaffected. The exception is presented in Sec. 9.3.5.

### 9.3. System tuning

This section discusses the experimental efforts to implement the protocol of optical pumping into a maximally entangled state. The first step of the experiment preparation was the calibration of the phase-modulated collective drive (A), which is discussed in Sec. 9.3.1. Subsequently, the first attempts at executing the protocol resulted in spectacularly low fidelities. While debugging the sequence, we discovered we were limited by spectator mode excitation, which created correlated bit-flip errors with high probability. After eliminating those (as discussed in Sec. 9.3.2), the protocol was still under-performing. We then discovered that while  $|\Psi^-\rangle$  is supposed to be the eigenstate of the protocol, there are in fact three experimental sources of direct coupling between  $|\Psi^-\rangle$  and  $|\Psi^+\rangle$ . Those are discussed in Sec. 9.3.3, 9.3.4 and 9.3.5. Unable to eliminate these couplings, we adapted the protocol to correct for them with a spin-echo technique discussed in Sec. 9.3.6. The adapted protocol allowed us to finally perform the experiment under favourable conditions and obtain the results presented in Sec. 9.4.

For consistency with Chapter 8, experiments described in this section use a dual notation. Experiments that only describe optical qubit operations refer to the qubit states as  $|1\rangle = |S_{1/2}, m_J = -1/2\rangle$  and  $|0\rangle = |D_{5/2}, m_J = -1/2\rangle$ . However, when discussing the collective optical pumping protocol, we refer to these as  $|\downarrow\rangle$  and  $|e\rangle$  respectively.

### 9.3.1. Phase-modulated drive


As discussed above in Sec. 9.1.3, the drive (A) is implemented as a sequence of two MS gates of length  $t_{\text{MS}}$  with identical spin phases, but with their motional phases offset by either  $\pi$  or  $(\pi - \delta t_{\text{MS}})$  (“Aussie style” as defined in Sec. 2.4.3).

Initially, we attempted to execute the drive as a single shaped pulse of length  $2 \times t_{\text{MS}}$  with an instantaneous phase update in the middle. However, this produced a large off-resonant carrier excitation. We later confirmed in simulations that indeed sudden phase changes cause the same kinds of errors as sudden Rabi frequency changes (Sec. 2.3.3). At the time, our hardware did not allow for smooth *phase* shaping, so we decided to separate the drive (A) into two individual pulses.

Each of the two pulses uses the same power, frequency and length settings, which are inherited from a calibrated MS gate (Sec. 8.2.4). However, the phases now require calibration due to a gap time  $\tau$  between the pulses. This can be understood as follows. For a single calibrated MS pulse, the spin phase  $\phi_s(t)$  evolves with time in the laboratory frame as


$$\phi_s(t) = \frac{1}{2}(\phi_b(t) + \phi(r)) = \omega_0 + \delta_{\text{com}}(t) \quad (9.16)$$

where  $\omega_0$  is the bare optical qubit frequency and  $\delta_{\text{com}}$  is the AC Stark shift during the MS pulse. This phase evolution is automatically tracked in hardware. However, during the gap time  $\tau$ , the qubit evolves at a bare frequency  $\omega_0$ , and accumulates an extra spin phase  $\delta_{\text{com}}\tau$  in the frame of the MS drive (to which the hardware is set). For the second pulse to be applied at the same spin phase as the first, it requires a phase offset of  $\Delta\phi_s \approx -\delta_{\text{com}}\tau$ , where the approximate sign applies as the AC Stark shift changes smoothly as the pulses are ramped on- and off. By a similar argument, the motional phase  $\phi_m$  requires a correction of  $\Delta\phi_m \approx -\delta\tau$ , where  $\delta$  is the gate sideband detuning.



**Figure 9.7:** Calibration of the two-loop phase-modulated drive (A). (Left) Calibration of the spin phase  $\phi_{s2}$  with two MS pulses of total length  $2 \times t_{\text{MS}}$  applied to an initial state  $|11\rangle$ . The correct phase corresponds to maximal  $P(00)$  (dashed line). (Right) Calibration of the motional phase  $\phi_{m2}$  with two MS pulses of total length  $\approx 1.5 \times t_{\text{MS}}$  applied to an initial state  $|11\rangle$ . The correct phase corresponds to minimum  $P(01 + 10)$  (dashed line).

These phases are calibrated as follows. First, the spin phase is calibrated by applying the MS drive of length  $2 \times t_{\text{MS}}$  to a pair of ions in  $|11\rangle$  and scanning the spin phase  $\phi_{s2}$  of the second pulse (Fig. 9.7 left). When  $\phi_{s2} = \phi_{s1}$ , the pulses pair maximises the probability of a full population transfer from  $|11\rangle$  to  $|00\rangle$ , as indicated as a dashed line in Fig. 9.7 (left), regardless of the motional phase setting. The signature of correct calibration of  $\phi_{m2}$  is maximum error resilience, which is difficult to probe at  $2 \times t_{\text{MS}}$ . Thus, we typically calibrate it by exploiting the fact that, for “Aussie-style” phase modulation, the phase-space loop is expected to close at all times, not just at  $2 \times t_{\text{MS}}$ . Fig. 9.7 (right) shows the result of scanning the phase offset  $\Delta\phi_m$  with  $\phi_{m2} = \phi_{m1} - \delta t + \Delta\phi_m$  and the total pulse length set to  $t \approx 1.5 \times t_{\text{MS}}$ . The dashed line indicates a phase offset with the lowest single-qubit excitation probability, indicating  $\alpha(t) \approx 0$  and hence a correct motional phase calibration.



**Figure 9.8:** (Left) Time trace of a calibrated Aussie-style phase modulated MS drive applied to an initial state  $|11\rangle$ . Full population transfer is achieved at time  $t = 2 \times t_{\text{MS}} \approx 150 \mu\text{s}$ . Trap potential glitches are noticeable as individual points which significantly differ from their immediate neighbours. (Right) Comparison between the odd populations measured at the end of the MS drive applied to an initial state  $|11\rangle$ .

The Fig. 9.8 (left) shows a typical time-trace of a calibrated phase-modulated (Aussie-style) MS drive. The odd population  $P(01 + 10)$  indeed remains close to 0 at all times, indicating successful phase correction, and its slow increase over time is caused by residual gate errors. A single cycle of the drive (A) corresponds to a pulse time of  $t = 2 \times t_{\text{MS}} \approx 150 \mu\text{s}$ . Fig. 9.8 (right) summarises typical errors of the drive. While the MS gate errors vary significantly over time (Sec. 8.3.3), the population error associated with a two-loop phase-modulated drive of length  $2 \times t_{\text{MS}}$  is approximately double that for a single-loop drive of length  $t_{\text{MS}}$ , and less than the error achievable without phase modulation<sup>6</sup>. Note that the final odd population  $P(01 + 10)$  measured with the initial state  $|11\rangle$  is exactly equal to the bit-flip probability  $p$  assuming a correlated error model (Sec. 9.2.2). While Fig. 9.8 (right) indicates  $p = 0.014(3)$  for a phase-modulated drive, this number varies over time as the MS gate errors drift, and can at times be as low as  $p = 0.010(1)$ .

<sup>6</sup> What escaped me at the time we ran the experiments was that, since the errors are proportional to the pulse time, it might be advantageous to apply drive (A) as a single loop of length  $\sqrt{2} \times t_{\text{MS}}$

### 9.3.2. Spectator modes

The first challenge of performing a high-fidelity dissipative pumping sequence was eliminating undesired spectator mode excitation. In order to avoid leakage out of  $|\Psi^-\rangle$ , the collective excitation drive must be sufficiently from any spectator transition out of either  $|\downarrow\rangle$  or  $|\uparrow\rangle$ . The difficulty of this process can be already explained with single-ion spectroscopy and is illustrated in Fig. 9.9. Low-power single-ion spectroscopy around the  $|\downarrow\rangle \leftrightarrow |e\rangle$  transition (Fig. 9.9 a)) reveals a small number of resonances associated with carrier and sideband transitions out of  $|\downarrow\rangle$  (left) and  $|\uparrow\rangle$  (right). To apply a low-error state-dependent force drive, the bichromatic laser field must be close to one of the resonances, and far detuned from any other transition out of any of the ground states. With a short low-power pulse, we find  $P(S_{1/2}) = 1$  for most laser frequencies, and it is therefore easy to ensure that the laser addresses one and only one transition.

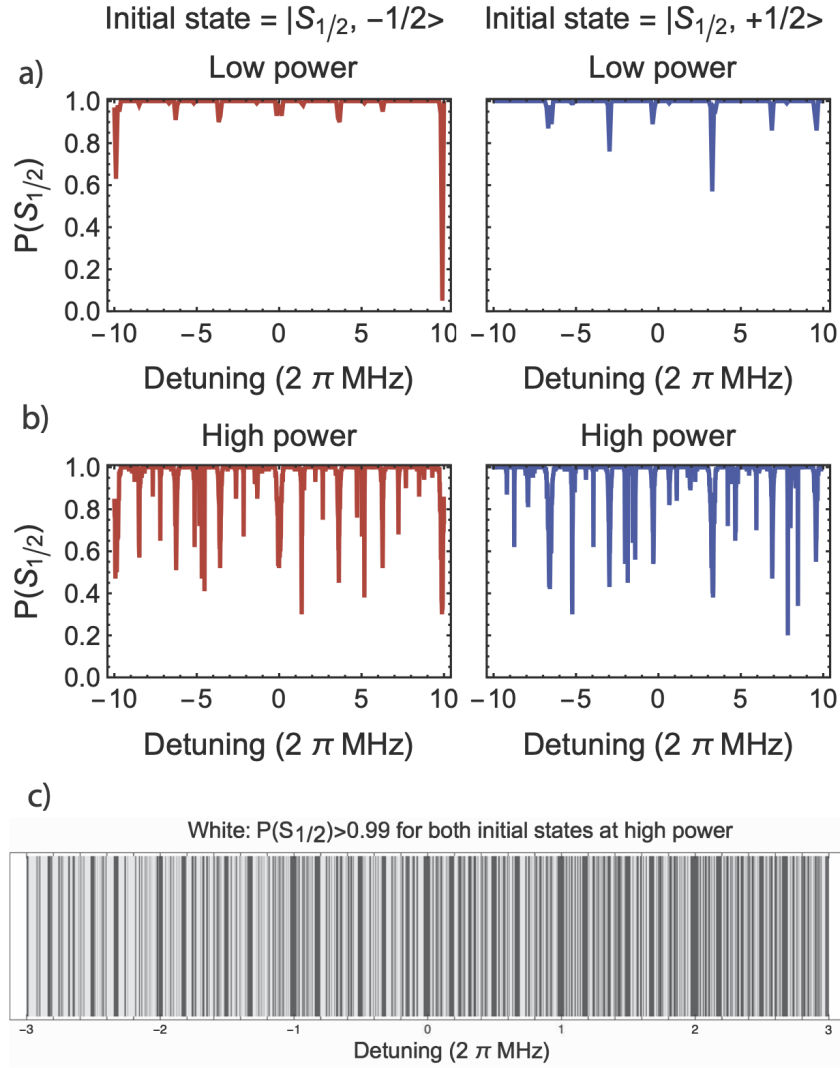
However, the MS drive is a long and high-power pulse. Performing single-ion spectroscopy with a shaped pulse of length  $t = 150\ \mu\text{s}$  and the same single-tone power as used for an MS gate, we find a spectrum shown in Fig. 9.9 b). Now, we find a large number of (sometimes overlapping) spectral features. In addition to the peaks seen in low-power spectroscopy, we find a large number of higher-order motional sidebands, as well as some possible micromotion sidebands of distant carrier transitions. It is now significantly more challenging to address one and only one resonance. With two ions, the spectrum is even more crowded due to an increase in the the number of motional modes.

This challenge is made doubly difficult by the fact even small single-ion excitation probability  $p$  becomes amplified by the dissipative protocol. For an example steady-state fidelity of 0.97, we require  $p < 0.01$ , which necessitates a detuning of over 5 linewidths from any resonance, not accounting for laser spectral impurity or mode frequency fluctuations. This is illustrated in Fig. 9.9 c), where the areas in white highlight the parts of the single-ion high-power spectrum where we measure  $p < 0.01$ .

Experimentally, we avoid spectator mode excitations during the drive (A) as follows. First, we calibrate a phase-modulated drive of length  $2 \times t_{\text{MS}}$  at a specific stretch-mode frequency  $\omega_{\text{str}}$  as described in Sec. 9.3.1. We then apply the drive for different “tilt curvatures” of the radial principal axes. The tilt potential is set to create confinement along  $+45^\circ$  deg to the trap normal, and anti-confinement along  $-45^\circ$  deg to the trap normal, leaving the axial potential unaffected. Thus, in the first approximation, changing the tilt curvature affects the radial mode frequencies and orientations (and hence, in turn, the radial mode temperatures), leaving the axial mode frequencies and the axial position unchanged.

For each setting of the tilt curvature, we measure the bit-flip probability in two ways. In the first, the ions are initialised in  $|\downarrow\downarrow\rangle$ , and we measure the probability  $p(\downarrow)$  for the drive (A) to cause a single-particle excitation<sup>7</sup>. In the second, the ions are initialised in  $|\uparrow\uparrow\rangle$  by means of optical pumping followed by a Zeeman pi-pulse, and we measure the probability  $p(\uparrow)$  for the drive (A) to cause a single-ion excitation. Both of these errors need to be


<sup>7</sup> In the optical qubit notation, this would be called  $P(01 + 10)$



**Figure 9.9:** Illustration of the ion spectral crowding challenge. a) Single-ion spectroscopy performed with short and low-power pulses. The ion's initial state is set to either  $|\uparrow\rangle$  (left column) or  $|\downarrow\rangle$  (right column). b) Spectral crowding observed with a single ion when the spectroscopy pulse has the same length and power as the drive (A) c) Illustration of the fine-grained nature of the spectral crowding when low excitation probability is required. Frequency is marked in white when the probability of excitation is below 0.01 for both initial states in the plots b), and in black otherwise.

## 9. OPTICAL PUMPING INTO A MAXIMALLY ENTANGLED STATE

simultaneously minimised, as the effective correlated bit-flip error  $p \approx (p(\uparrow) + p(\downarrow))/2$ . If no tilt setting creates sufficiently low  $p$ , we change  $\omega_{\text{str}}$  and repeat the procedure.



**Figure 9.10:** Bit-flip errors  $p(\downarrow)$  and  $p(\uparrow)$  measured for different settings of the radial mode tilt curvature. At the tilt curvature is increased, the lower radial centre-of-mass mode decreases in frequency from  $\approx 2\pi \times 4$  MHz to  $\approx 2\pi \times 3.5$  MHz. At the same time, the upper radial centre-of-mass mode increases in frequency from  $\approx 2\pi \times 5$  MHz to  $\approx 2\pi \times 5.4$  MHz. Larger tilt curvatures also correspond to mode orientations closer to  $\pm 45^\circ$  from the trap normal/horizontal. Following a low-resolution scan over a broad region of settings (left), the plot on the right shows a precise measurement within the promising region (black rectangle). Following the scan, the tilt is set to the low-error region (green rectangle). Individual measurements are connected with straight lines only to guide the eye.


Fig. 9.10 illustrates the calibration procedure at  $\omega_{\text{str}} = 2\pi \times 2.4$  MHz. We apply a tilt curvature in the range of approximately  $(3 - 11) \times 10^7$  V/m<sup>2</sup>, which corresponds to a frequency span of  $\approx 2\pi \times (4 - 3.5)$  MHz and  $\approx 2\pi \times (5 - 5.4)$  MHz for the lower and the upper radial COM mode respectively. Higher tilt curvatures correspond larger radial mode splittings. A broad scan of the radial mode tilt (Fig. 9.10 left) reveals a narrow region of settings (black rectangle) where both  $p(\downarrow)$  and  $p(\uparrow)$  are low. We then perform a precise scan in that narrow region (Fig. 9.10 right), finding the optimal tilt setting of  $\approx 5.9 \times 10^7$  V/m<sup>2</sup> (green rectangle). Typically, with the tilt optimised, we find that  $p(\downarrow) \approx 0.01$ , while  $p(\uparrow)$  is usually a few times lower, on the order of  $\approx 0.003$ . This is expected, as  $p(\uparrow)$  is limited only by spectator mode excitation, while  $p(\downarrow)$  is increased by both spectator mode coupling and residual spin-motion entanglement of the drive (A).

Off-resonant mode excitation is one of the more challenging aspects of the experimental implementation of the protocol. The radial mode tilt calibration, while largely reproducible, is slow and needs to be repeated frequently, especially after any significant trap potential drift events (e.g. ion loading or changing the laser duty cycle). While we usually assume the error to be dominated by  $p(\downarrow)$ , sometimes  $p(\uparrow)$  becomes significant, complicating the error analysis. Furthermore, it is difficult to construct an *a priori* error model, since it is challenging to estimate to what extent  $p(\downarrow)$  is limited by the imperfections in the two-level Hamiltonian, and to what extent by spectator mode couplings.

Future implementations of the collective optical pumping protocol should anticipate this error and take precautions at the experiment design stage. The density of spectator transitions can be reduced in several ways. While our trap-integrated laser at  $\lambda = 729$  nm has significant overlap with all motional modes, aligning it orthogonal to a subset of modes will reduce spectral crowding. Furthermore, while our EIT cooling only addresses the axial modes (Sec. 5.6.3), a change in beam alignment could allow for multi-mode cooling of all axial and radial modes. Alternatively, a different multi-mode cooling technique could be used to cool all the modes with a significant Lamb-Dicke parameter close to the motional ground state [Jos20]. This should decrease the number of spectator modes, and reduce coupling to higher-order motional sidebands. Finally, higher B-fields result in a larger spectral separation between different carrier transitions, in turn reducing the spectral crowding.

### 9.3.3. Effective potential anharmonicity

The most striking difference between the laser-ion interactions observed in the trap #4 and #5 were the sideband oscillations. While the BSB trace of a ground-state cooled axial stretch mode in the trap #4 was almost identical to the theoretical curve (Sec. 8.1.2), the same experiment in the trap #5 revealed a large imbalance of the sideband Rabi frequencies in the working-well, as shown in Fig. 9.11 (left).



**Figure 9.11:** BSB oscillations on the ground-state cooled axial stretch mode in the trap #5, together with theoretical traces (solid lines) that include a Rabi frequency imbalance and finite temperature (Left) BSB oscillations in the working-well. The rapid increase of  $P(01 + 10)$  indicate effective anharmonicity (Right) BSB oscillations in the well with an additional cubic term added to the potential. The increase in  $P(01 + 10)$  is now slower, indicating a decrease in the effective anharmonicity.

The observed rise in the moving average of  $P(01 + 10)$  indicates that the two ions experience different sideband Rabi frequencies. Since the carrier Rabi frequencies are still equalised to within 1%, this must be caused by a difference of the Lamb-Dicke parameters.

## 9. OPTICAL PUMPING INTO A MAXIMALLY ENTANGLED STATE

The theoretical trace in Fig. 9.11 (left) corresponds to

$$|1 - \eta_1/\eta_2| = 0.1,$$

where  $\eta_1$  and  $\eta_2$  are the Lamb-Dicke parameters of ion 1 and 2 respectively.

When we first observed this effect, we conjectured that the difference in  $\eta$  is caused by trap anharmonicity, which changes the effective trap frequency already on the scale of  $\sim 1 \mu\text{m}$ . Indeed, by adding additional cubic terms to the working-well potential, we could reduce this imbalance to  $|1 - \eta_1/\eta_2| \approx 0.06$  (Fig. 9.11 right), beyond which we became limited by the amplifiers. However, the calibration of the cubic term was gruesome due to a cross-talk between the cubic, quadratic and linear terms, and we did not use it in other experiments described in this chapter.

We conjectured that the strong anharmonicity must be caused by charges deposited on the dielectrics nearby the ions. What did not occur to us at the time was that gradients in  $\eta$  could also be caused by a complex tightly-focused beam profile (Sec. 4.6.3). Indeed, simulations presented in the PhD thesis of C. Zhang suggest that the changes of the effective laser  $k$ -vector along the trap axis could account for the significant fraction or even all of the observed imbalance in  $\eta$  [Zha21]. This hypothesis could have been easily verified by driving the BSB with the free-space  $\lambda = 729 \text{ nm}$  “debug” beam, but unfortunately, we did not do that at the time. I will henceforth remain agnostic about the source of the differential sideband Rabi frequency, and refer to the effect as “effective anharmonicity”.

The effective anharmonicity does not increase the MS gate error, but causes coupling between  $|\Psi^-\rangle$  and  $|\Psi^+\rangle$  as follows. Consider the evolution of the state  $|\psi(0)\rangle = |\downarrow\uparrow\rangle$  during the drive (A). Since only one of the particles participates in the interaction, the system evolution is that of a single-ion state-dependent force acting on the qubit in  $|\downarrow\rangle$  (Sec. 2.4.1), with the qubit in  $|\uparrow\rangle$  unaffected. Assuming  $\alpha(t) = 0$ , the result from Eq. 2.31 implies that the state of the ions at time  $t$  is given by:

$$|\psi(t)\rangle = e^{-i\Phi_1(t)}|\downarrow\uparrow\rangle \quad \text{where} \quad \Phi_1(t) = \frac{\eta_1^2 \Omega^2 m \pi}{2\delta^2}$$

and  $\eta_1$  is the Lamb-Dicke parameter of ion 1. Likewise, for two ions prepared at  $t = 0$  in  $|\psi(0)\rangle = |\uparrow\downarrow\rangle$  we find

$$|\psi(t)\rangle = e^{-i\Phi_2(t)}|\uparrow\downarrow\rangle \quad \text{where} \quad \Phi_2(t) = \frac{\eta_2^2 \Omega^2 m \pi}{2\delta^2}$$

where  $\eta_2$  is the Lamb-Dicke parameter of ion 2. Now consider an anharmonicity parameter  $\epsilon$  such that  $\eta_1 = \eta(1 - \epsilon/2)$  and  $\eta_2 = \eta(1 + \epsilon/2)$ . For  $\epsilon \ll 1$ , the evolution can be written as

$$\begin{aligned} |\downarrow\uparrow\rangle &\longrightarrow e^{-i(1-\epsilon)\Phi(t)}|\downarrow\uparrow\rangle \\ |\uparrow\downarrow\rangle &\longrightarrow e^{-i(1+\epsilon)\Phi(t)}|\uparrow\downarrow\rangle \end{aligned}$$



where  $\Phi(t) = \frac{\eta^2 \Omega^2 m \pi}{2 \delta^2}$ . In other words, the system evolves as if these eigenstates were non-degenerate.

Now suppose that the system is prepared in  $|\Psi^-\rangle = (|\downarrow\uparrow\rangle - |\uparrow\downarrow\rangle)/\sqrt{2}$ . After a single cycle of the drive (A), it evolves into

$$\begin{aligned} |\Psi^-\rangle &\longrightarrow \frac{1}{\sqrt{2}}(e^{-i(1-\epsilon)\Phi(t)}|\downarrow\uparrow\rangle - e^{-i(1+\epsilon)\Phi(t)}|\uparrow\downarrow\rangle) \\ &= \cos(\Delta\phi/2)|\Psi^-\rangle + i\sin(\Delta\phi/2)|\Psi^+\rangle \end{aligned}$$

where  $\Delta\phi = 2\epsilon\beta(t)$ . Using  $\epsilon = 0.1$  as found above and  $\Phi(t) = \pi/4$  which is the optimal setting of the drive (A), we find  $\Delta\phi = 9$  deg and a corresponding probability of leakage out  $|\Psi^-\rangle$  of given by

$$p = \sin^2(\Delta\phi/2) = 6 \times 10^{-3} \text{ per cycle.}$$

### 9.3.4. B-field gradient

The second source of coupling between  $|\Psi^-\rangle$  and  $|\Psi^+\rangle$  are the B-field gradients. Due to imperfect alignment between the B-field “coils” centre and the trap centre, we expect a certain axial gradient  $\partial|B|/\partial x$  of the total B-field magnitude (Sec. 3.4.3). This lifts the degeneracy between  $|\downarrow\uparrow\rangle$  and  $|\uparrow\downarrow\rangle$ , causing a frequency difference of

$$\Delta\omega = g_J\mu_B \frac{\partial|B|}{\partial x} \Delta x + O(\Delta x^2)$$

where  $g_J = 2$  is the Lande factor for the ground state manifold,  $\mu_B$  is the Bohr magneton and  $\Delta x$  is the ion spacing. As a result, a system prepared in  $|\Psi^-\rangle$  at time  $t = 0$  evolves into

$$|\Psi^-\rangle \longrightarrow \cos(\Delta\phi/2)|\Psi^-\rangle + i\sin(\Delta\phi/2)|\Psi^+\rangle, \quad (9.17)$$

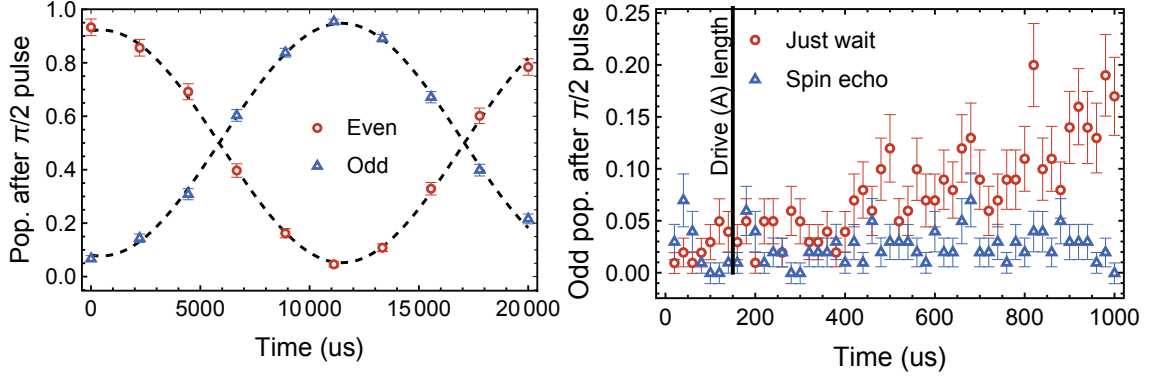
where

$$\Delta\phi = \Delta\omega t.$$

This coupling can be observed experimentally as follows. We prepare the ions in  $|\downarrow\downarrow\rangle$  and perform an MS gate on the optical qubit transition, followed by a  $\pi/2$  “analysis” pulse with its phase calibrated to create  $(|\downarrow e\rangle + |e \uparrow\rangle)/\sqrt{2}$ . We then apply a carrier “shelve”  $\pi$ -pulse on the  $|\uparrow\rangle \leftrightarrow |e\rangle$  transition, which prepares  $|\Psi^+\rangle$ . We let the system evolve freely for time  $t$ , during which it transforms into a superposition of  $|\Psi^+\rangle$  and  $|\Psi^-\rangle$  as described by Eq. (9.17). Afterwards, we perform a  $\pi/2$  “analysis” pulse on the Zeeman qubit, which maps  $|\Psi^+\rangle$  into  $(|\downarrow\downarrow\rangle + e^{i\phi}|\uparrow\uparrow\rangle)/\sqrt{2}$ , but leaves  $|\Psi^-\rangle$  unaffected. We then perform Zeeman qubit readout (Sec. 2.3.7). Any of the “even” outcomes (00 or 11) implies the ions were in  $|\Psi^+\rangle$  before the analysis pulse, while “odd” outcomes correspond to  $|\Psi^-\rangle$ .

Fig 9.12 (left) shows the observed coupling between  $|\Psi^+\rangle$  and  $|\Psi^-\rangle$  in absence of any

## 9. OPTICAL PUMPING INTO A MAXIMALLY ENTANGLED STATE



**Figure 9.12:** (Left) Oscillations between the odd and even populations after an analysis that reveal a coupling between  $|\Psi^+\rangle$  and  $|\Psi^-\rangle$ . Dashed lines show a sinusoidal fit (Eq. (9.17)). (Right) The change in the  $|\Psi^+\rangle$  population for short wait times with and without spin echo. For  $t \approx 165 \mu\text{s}$  (corresponding to a single cycle of the protocol), we record negligible population transfer in either case.

control fields at the axial centre-of-mass frequency of  $\omega = 2\pi \times 1.37 \text{ MHz}$  (ion spacing  $\Delta x \approx 4.5 \mu\text{m}$ ). The population exchange rate fit implies a frequency splitting of

$$|\Delta\omega| = 2\pi \times 45.4(10) \text{ Hz}$$

corresponding to a B-field gradient of

$$\frac{\partial|B|}{\partial x} = 3.6(1) \text{ G/m}$$

consistent with the expected coil misalignment. We also verified that increasing the ion spacing  $\Delta x$  increases the frequency splitting  $\Delta\omega$ .


While the gradient could be in principle decreased by better coil alignment, the singlet-triplet coupling is slow enough to contribute little to the protocol error. For a single protocol cycle of length  $t \approx 165 \mu\text{s}$ , we find  $\Delta\phi = 2.7 \text{ deg}$  and the corresponding leakage probability of

$$p = \sin^2(\Delta\phi/2) = 6 \times 10^{-4} \text{ per cycle.}$$

We attempted to measure  $p$  experimentally for short evolution times  $t$  in Fig. 9.12 (right). While the noise in  $|\Psi^+\rangle$  preparation probability does allow to precisely measure  $p$  at  $t \approx 165 \mu\text{s}$ , we do find that the error can be effectively cancelled by adding a Zeeman spin-echo pi-pulse in the middle of the free evolution. The adapted protocol discussed in Sec. 9.3.6 includes this spin-echo, and thus the error associated with the B-field gradient is expected to be negligible.

### 9.3.5. Differential AC Stark shift

Combining the effects of the effective anharmonicity and the B-field gradient, we expected the drive (A) to cause a population exchange between  $|\Psi^+\rangle$  and  $|\Psi^-\rangle$  with probability  $p \lesssim 7 \times 10^{-3}$  at  $t_A = 150 \mu\text{s}$ . However, upon repeating the singlet-triplet coupling rate measurement from Sec. 9.3.4 with the drive (A) applied during the wait time, we found a strikingly larger coupling rate, as shown in Fig. 9.13.



**Figure 9.13:** The measured population exchange between  $|\Psi^+\rangle$  and  $|\Psi^-\rangle$  when the drive (A) is applied. Dashed lines indicate a sinusoidal fit with a Gaussian decay. The coupling rate significantly exceeds the expectation given the observed anharmonicity and the B-field gradient.

The fit in Fig. 9.13 reveals that, during an MS interaction, the states  $|\downarrow\uparrow\rangle$  and  $|\uparrow\downarrow\rangle$  experience an effective frequency splitting of

$$|\Delta\omega| = 2\pi \times 2.13(3) \text{ kHz}.$$


For a single application of the drive (A) with length  $t = 155 \mu\text{s}$ , this causes a phase shift of  $\Delta\phi = 119(2) \text{ deg}$  and a corresponding bit-flip error of

$$p = 0.74(1),$$

orders of magnitude higher than what would be expected given the anharmonicity and the B-field gradients. Indeed, even after the problem with the spectator mode off-resonant excitation was solved, this effect alone led to a complete failure of the protocol.

In investigating the origin of this coupling we found, to our surprise, that the degree of carrier Rabi frequency imbalance differs significantly between different optical qubit transitions. Most importantly, when the fractional carrier Rabi frequency imbalance on the gate transition  $|\downarrow\rangle \leftrightarrow |e\rangle$  is  $\leq 0.01$  (Sec. 8.1.3), we find a Rabi frequency difference of  $\epsilon \approx 5\%$  on a carrier transition  $|\uparrow\rangle \leftrightarrow |D_{5/2}, m_J = +3/2\rangle$ , which is detuned by  $\approx 2\pi \times 0.9 \text{ MHz}$  from the BSB pulse of the MS Hamiltonian. This is illustrated in Fig. 9.14.

## 9. OPTICAL PUMPING INTO A MAXIMALLY ENTANGLED STATE




**Figure 9.14:** Measurements of the optical qubit carrier Rabi oscillations at the working-well. Solid lines show theoretical traces that include a decay rate (identical to both traces) and a Rabi frequency difference (different between the two traces). (Left) Carrier oscillations on the  $|\downarrow\rangle \leftrightarrow |e\rangle$  transition. The fit indicates a Rabi frequency imbalance  $\epsilon < 0.01$ , as the working-well was calibrated to equalise the single-ion Rabi frequencies (Right) Carrier oscillations on the  $|\uparrow\rangle \leftrightarrow |D_{5/2}, m_J = +3/2\rangle$  transitions. The fit indicates a Rabi frequency imbalance  $\epsilon = 0.05$ . Other carrier transitions exhibit different degrees of Rabi frequency difference.

The observed singlet-triplet coupling can then be explained as follows. The BSB drive of the MS Hamiltonian changes the Zeeman qubit frequency by  $\Delta\omega' \approx 2\pi \times 20$  kHz due to an AC Stark shift of the  $|\uparrow\rangle \rightarrow |D_{5/2}, m_J = +3/2\rangle$  spectator transition. The Rabi frequency imbalance  $\epsilon$  on that spectator transition leads to a frequency splitting of  $|\Delta\omega| = 2\epsilon\omega' \approx 2\pi \times 2$  kHz of the nominally degenerate states  $|\downarrow\uparrow\rangle$  and  $|\uparrow\downarrow\rangle$ , causing the observed singlet-triplet coupling.

As expected from this theory, we found that large  $|\Delta\omega|$  could be induced already with a single tone, and that it could be amplified by decreasing the detuning  $\delta$  from the  $|\uparrow\rangle \rightarrow |D_{5/2}, m_J = +3/2\rangle$  transition. Fig. 9.15 (red data) shows the result of preparing  $|\Psi^+\rangle$  and then applying a single pulse with detuning  $\delta \approx -2\pi \times 0.35$  MHz (and off-resonant from all other transitions) and power comparable to that used in the MS interaction. This creates a frequency splitting of  $|\omega| \approx 2\pi \times 6$  kHz. In the subsequent experiments, we also used these pulse settings as a Zeeman qubit single-ion addressing drive. This allowed us to prepare states with high overlap with  $|\Psi^-\rangle$  without translating the ions across the tightly-focused  $\lambda = 729$  nm beam.

We still do not fully understand all of these observations. First, while the observed Rabi frequency imbalance  $\epsilon$  accounts for the measured singlet-triplet coupling rate, it is not fully clear why  $\epsilon$  is different for different carrier transitions in the first place. Second, the fits in Fig. 9.13 and Fig. 9.15 exhibit strong contrast decay. Assuming a Gaussian noise model, we find a fractional standard deviation on  $\Delta\omega$  equal to

$$\sigma_{\Delta\omega}/\Delta\omega = 0.18(2).$$



**Figure 9.15:** The measured population exchange between  $|\Psi^+\rangle$  and  $|\Psi^-\rangle$  with a single-tone  $\lambda = 729$  nm pulse near-detuned to the  $|\uparrow\rangle \rightarrow |D_{5/2}, m_J = +3/2\rangle$  transition. The dashed line shows a sinusoidal fit with a Gaussian decay. The fitted decay rate indicates differential Rabi frequency noise which significantly exceeds the noise measured on any of the carrier transitions. Blue points show the population exchange with a Zeeman spin-echo pi-pulse inserted in the middle of the  $\lambda = 729$  nm pulse. At the length of the drive (A)  $t = 150$   $\mu$ s (solid line), we record no population transfer with the spin echo, even though the triplet-singlet coupling rate exceeds that recorded for the drive (A).

To account for this noise solely by the AC Stark shift discussed above with  $\epsilon$  and  $\delta$  held constant, we would require a fractional error in the Rabi frequency of

$$\sigma_\Omega / \langle \Omega \rangle = 0.09(1).$$

However, that is incompatible with the fits to the two-ion carrier oscillations, which indicate a noise level of  $\sigma_\Omega / \langle \Omega \rangle \approx 0.015$ , comparable for different carrier transitions (Fig. 9.14)<sup>8</sup>. This indicates that the decay is dominated by fluctuations in the Rabi frequency *difference*, not the average Rabi frequency! Nonetheless, introducing a Zeeman spin-echo pi-pulse in the middle of the sequence effectively eliminates any observable singlet-triplet coupling for  $t \leq 200$   $\mu$ s (Fig 9.15), suggesting that the fluctuations in the Rabi frequency difference are predominantly slow.

We have still not reached a satisfactory theory that explains both of these puzzles. Our current best guess is that differences in  $\epsilon$  across different carrier transitions are caused by secondary polarization components in the beam profile away from the beam centre (Sec. 4.6.3). Detailed simulations of a single ion interacting with a tightly-focused laser beam described in the PhD thesis of C. Zhang reveal that the Rabi frequencies of different carrier transitions are indeed maximised at different locations [Zha21]. It is possible that these simulations can reproduce the observed gradients, but this work is still in progress

<sup>8</sup> In a separate experiment, we verified that the fractional noise in the total AC Stark shift experienced by two ions in the ground state is  $\approx 0.03$ , consistent with the carrier oscillation fits and significantly below the noise on the differential AC Stark shift

at the time of writing. The increased noise in the Rabi frequency difference could then be explained by the ion sampling a region of space where the Rabi frequency gradient depends strongly on the position. The sampling could be due to thermal wavepacket extent, residual excess micromotion or fast stray field fluctuations.

### 9.3.6. Adapted protocol

As discussed in Sec. 9.3.5, the drive (A), while nominally leaving the singlet state  $|\Psi^-\rangle$  unaffected, actually couples it to  $|\Psi^+\rangle$  with probability  $p = 0.74$  per cycle. Fortunately, this coupling appears to be coherent and reversible with a spin-echo pulse.

We can introduce a spin-echo element to our protocol as follows. Instead of ending every cycle with a symmetric ground-state drive (C) of optimal area  $\theta \approx 0.72\pi$ , we alternate between two pulse area settings. Every odd cycle, we use  $\theta = \pi$ , which implements a spin echo. Every even cycle, we apply  $\theta = \pi/2$ , which we calculated to be the optimal setting. Simulations confirm that this modification cancels out the steady-state error caused by any offsets in  $\Delta\omega$  which are constant within a pair of cycles. At the same time, the adapted protocol converges at only a slightly slower rate of  $N_0 = 8.0$  cycles, assuming the branching ratio is still set to the optimal value of  $\gamma \approx 0.22\pi$ .

Experimentally, we verified that the singlet leakage probability reduces to  $p \lesssim 0.01$  in the adapted protocol. Below this number, it is difficult to distinguish this leakage channel from other errors present in the protocol. Note that the spin-echo is effective at returning the population from  $|\Psi^+\rangle$  to  $|\Psi^-\rangle$  despite a repump process that changes the total population in both states.

## 9.4. Experiment and results

In this section, we experimentally demonstrate optical pumping into a maximally entangled state using the adapted protocol described in Sec. 9.3.6.

### 9.4.1. Experimental implementation

The data acquisition period begins by fixing the average power of the  $\lambda = 729$  nm laser to reduce trap potential fluctuations. After the mode frequencies are stabilised and calibrated, the protocol operation can begin. Every experiment begins by Doppler cooling all motional modes, EIT-cooling the axial modes and sideband-cooling the axial stretch mode (Sec. 8.1.2). We then prepare different possible initial states of the ions. Following internal-state preparation, we apply  $N$  individual cycles of the protocol. Afterwards, we measure the fidelity of  $|\Psi^-\rangle$ , which requires only three separate measurement settings. In this section, we explain in depth each of those steps.

### Laser power stabilisation

As discussed in Sec. 8.2.2, laser-induced frequency jumps can be eliminated by stabilising the average optical power through trap integrated waveguides. This is particularly important in the dissipative protocol, where changes in trap frequencies can lead to spectator mode excitation (Sec. 9.3.2). However, unlike the MS gate, the dissipative sequence is not fixed in length, and increasing the number of cycles should in principle lead to continual fidelity improvements. On the other hand, longer sequences correspond to larger average  $\lambda = 729$  nm power, causing drifts.

In order to make the data acquisition practical, we set the maximum length of high-power  $\lambda = 729$  nm light per experimental shot to 3 ms, limiting the maximum number of protocol cycles to  $N = 16$ . This was measured to be sufficient to achieve a steady-state for a large range of input states, and corresponds to  $\approx 1/e^2$  time of the protocol. In order to avoid frequency jumps, every calibration experiment running in the background also includes a 3 ms laser pulse at the same power. This slows down the protocol calibration and data acquisition, but still maintains a manageable level. On the other hand, increasing the number of cycles by an order of magnitude would result in impractically slow calibrations given the 1 – 2 hour lifetime of a two-ion crystal recorded in trap #5.

### Internal-state preparation

While the ideal dissipative protocol creates the same steady-state regardless of the input state, we do explicitly initialise the input state for two reasons. First, in the initial experiments, we vary the input state to verify whether it has any effect on the protocol outcome. Second, in the later experiments, we prepare the system with initial singlet fidelity of  $F(|\Psi^-\rangle) \approx 70\%$ . This allows us to guarantee that the protocol approximates its steady-state within  $N = 16$  cycles (Sec. 9.4.3).

In order to prepare the initial state with a given  $F(|\Psi^-\rangle)$ , we first prepare  $|\Psi^+\rangle$  as described in Sec. 9.3.4. We then apply a single-ion addressing drive which exploits the observed differential AC Stark shift to map  $|\Psi^+\rangle$  into  $|\Psi^-\rangle$ . We can control the initial singlet fidelity either by varying the length of the single-ion addressing drive (which affects the relative populations in  $|\Psi^\pm\rangle$ ) or by changing the phase of the  $\pi/2$  analysis pulse used to prepare  $|\Psi^+\rangle$ . The latter method redistributes some of the population into  $|\downarrow\downarrow\rangle$  and  $|\uparrow\uparrow\rangle$  as well. The fidelity of this unitary preparation of  $|\Psi^-\rangle$  is limited to  $\approx 88\%$  due to noise in the differential AC Stark shift (Sec. 9.3.5).

### Single cycle

Each cycle of the protocol proceeds as follows. First, we apply the drive (A) using a phase-modulated MS sequence as described in Sec. 9.3.1, with a total duration of  $t_A \approx 150$   $\mu$ s and  $\Phi \approx \pi/4$ . Then, the state  $|e\rangle$  is repumped through a simultaneous application of  $\lambda = 854$  nm and  $\lambda = 866$  nm drives for  $t_B \approx 5$   $\mu$ s. For this decay, we measure a branching

## 9. OPTICAL PUMPING INTO A MAXIMALLY ENTANGLED STATE

ratio of  $p_{e \rightarrow \downarrow} \approx \frac{2}{3}$  and  $p_{e \rightarrow \uparrow} \approx \frac{1}{3}$ , corresponding to

$$\gamma \approx 0.3\pi.$$

This branching ratio can be understood as follows. Since the  $\lambda = 854 \text{ nm}$  laser is primarily  $\pi$ -polarized, the dominant excitation channel is  $|e\rangle \rightarrow |P_{3/2}, m_j = -1/2\rangle$ . This  $P_{3/2}$  state decays with branching ratios  $2/3$  and  $1/3$  into  $|\downarrow\rangle$  and  $|\uparrow\rangle$  specified by the Clebsch-Gordon coefficients (App. 11.1). After  $5 \mu\text{s}$ , we measure a repump probability of  $> 0.9999$ , and no detectable loss outside the  $\{S_{1/2}, D_{5/2}\}$  subspace. We sometimes record transient population trapping in  $|D_{5/2}, m_J = -5/2\rangle$ , which is not efficiently repumped by the predominantly  $\pi$ -polarized  $\lambda = 854 \text{ nm}$  laser. However, we typically find  $< 10^{-4}$  of the population in  $|D_{5/2}, m_J = -5/2\rangle$  at any given moment, making the repumping error negligible.

The symmetric drive (C) is implemented identically to single Zeeman qubit rotations described in Sec. 5.5.3. To implement the adapted protocol, we alternate the drive time between  $t_C = 6.4 \mu\text{s}$  ( $\theta_C = \pi$ ) in odd cycles and  $t_C = 3.2 \mu\text{s}$  ( $\theta_C = \pi/2$ ) in even cycles. Because the distance between the ions and the PCB track through which the current is passed ( $\approx 1 \text{ mm}$ ) vastly exceeds the ion spacing ( $\approx 5 \mu\text{m}$ ), we expect a negligible difference between the single-ion Rabi frequencies  $\Omega_C^{(1)}$  and  $\Omega_C^{(2)}$ . Indeed, we did not manage to measure any differential Rabi frequency even for long drive times, allowing us to comfortably constraint the the difference to  $|1 - \Omega_C^{(2)}/\Omega_C^{(1)}| \ll 0.01$ . Therefore, Zeeman Rabi frequency imbalance contributes  $\epsilon \ll 10^{-4}$  to the protocol error (Sec. 9.2.3). Each protocol cycle takes  $\approx 165 \mu\text{s}$ , which includes gap times between pulses.

### Fidelity measurement

As discussed in Sec. 2.4.4, Bell-state fidelities can be extracted from only three parity measurements  $\langle \sigma_z \sigma_z \rangle$ ,  $\langle \sigma_x \sigma_x \rangle$  and  $\langle \sigma_y \sigma_y \rangle$ . Interestingly, measuring the fidelity of the odd Bell states  $|\Psi^\pm\rangle$  is in practice significantly easier than for even Bell states  $|\Phi^\pm\rangle$ . This is because, while the latter demonstrate parity oscillations as a function of the analysis pulse phase, the  $x/y$  parity is expected to be constant for the former ( $\langle \sigma_x \sigma_x \rangle = \langle \sigma_y \sigma_y \rangle$ ).

This leads to a simple fidelity measurement method. Consider a parity measurement after a global  $R(\pi/2, \phi) \otimes R(\pi/2, \phi)$  “analysis” pulse with an outcome  $\langle \sigma_\phi \sigma_\phi \rangle$ , and another parity measurement after a global  $R(\pi/2, \phi + \pi/2) \otimes R(\pi/2, \phi + \pi/2)$  “analysis” pulse with an outcome  $\langle \sigma_{\phi+\pi/2} \sigma_{\phi+\pi/2} \rangle$ . By direct calculation, it is possible to show that the sum of these two average outcomes is independent of  $\phi$  and therefore

$$\langle \sigma_x \sigma_x \rangle + \langle \sigma_y \sigma_y \rangle = \langle \sigma_\phi \sigma_\phi \rangle + \langle \sigma_{\phi+\pi/2} \sigma_{\phi+\pi/2} \rangle.$$

In other words, in order to estimate the fidelity of  $|\Psi^\pm\rangle$ , in addition to measuring  $\langle \sigma_z \sigma_z \rangle$ , it is sufficient to measure the parity following two analyses pulses with phases offset by  $\pi/2$ , regardless of the actual pulse phase<sup>9</sup>.

<sup>9</sup> Alternatively, the analysis pulse can be applied with its phase randomised shot-to-shot. This is because



Thanks to this insight, we measure the fidelities as follows. Following  $N$  cycles of the collective optical pumping protocol, we either apply no analysis pulse (setting 1), or apply a global Zeeman qubit  $\pi/2$ -rotation with the phase set to either  $\phi = 0$  (setting 2) or  $\phi = \pi/2$  (setting 3). Afterwards, we perform two-ion Zeeman qubit readout, shelving the populations into one or multiple  $D_{5/2}$  sublevels. We correct for readout and shelving errors (Sec. 8.1.4) and extract the probabilities  $P(\downarrow\downarrow)$ ,  $P(\downarrow\uparrow + \uparrow\downarrow)$  and  $P(\uparrow\uparrow)$  corresponding to zero, one, and two ions in  $|\downarrow\rangle$  respectively. For each setting, we estimate the parity as  $\langle\sigma_i\sigma_i\rangle = P(\downarrow\downarrow) + P(\uparrow\uparrow) - P(\downarrow\uparrow + \uparrow\downarrow)$ , where  $i = z$  for setting 1,  $i = x$  for setting 2 and  $i = y$  for setting 3. We then estimate the fidelities  $F(|\Psi^\pm\rangle)$  using Eq. (2.34).

### 9.4.2. Dynamical behaviour

Following the protocol turn-on, we begin to observe convergence dynamics superimposed with strong population oscillations between  $|\Psi^-\rangle$  and  $|\Psi^+\rangle$ , as illustrated in Fig 9.16a. The oscillations are indeed expected given the differential AC Stark shift (Sec. 9.3.5), and high-fidelity singlet states are only produced after even cycles in the adapted protocol (Sec. 9.3.6). Since figures such as Fig 9.16a are visually hard to interpret, from now on I only measure and plot the fidelities after even cycles.

To test the protocol, we vary the length of the single-ion addressing pulse to initialise the ions in a superposition of  $|\Psi^-\rangle$  and  $|\Psi^+\rangle$ . Then we apply up to  $N = 16$  cycles of the protocol. Fig. 9.16b shows how all the input states converge towards  $|\Psi^-\rangle$ . In this dataset, all the input states converge towards the same final state with  $F(|\Psi^-\rangle) \approx 0.9$ . The states with lower initial fidelity have clearly not yet reached their steady-state at  $N = 16$ . This is consistent with the convergence rate of  $N_0 \approx 8$  cycles, where for an ideal error-free protocol the state with  $F(|\Psi^-\rangle) = 0$  at  $N = 0$  is expected to have a fidelity of  $F(|\Psi^-\rangle) = 1 - 1/e^2 \approx 0.86$  at  $N = 16$ . On the other hand, the input state with the largest initial fidelity  $F(|\Psi^-\rangle) \approx 0.77$  (yellow trace in Fig. 9.16b) is already at the quasi steady-state at  $N = 16$ .

### 9.4.3. Steady-state fidelity


The dynamical measurements revealed that high singlet fidelity at  $N = 16$ , which is our proxy for the steady-state fidelity, requires high singlet fidelity at  $N = 0$ . Therefore, to estimate the steady-state fidelity, we unitarily prepare  $|\Psi^-\rangle$  with close-to-maximum fidelity, and compare it to the fidelity after  $N = 16$  cycles.

Fig. 9.17 shows the resulting input-output relation. We vary the initial fidelity by adding a phase offset to the unitary preparation pulse, thus changing the population in  $|\downarrow\downarrow\rangle$  and  $|\uparrow\uparrow\rangle$ . This changes the input singlet fidelity (green points and line) between a maximum value of  $F(|\Psi^-\rangle) \approx 0.88$  and  $F(|\Psi^-\rangle) \approx 0.72$  for a  $\pm 25$  deg phase offset. At the same time, the singlet fidelity at the protocol output (blue points) is approximately independent of


---

$\langle\sigma_x\sigma_x\rangle + \langle\sigma_y\sigma_y\rangle = 2\langle\sigma_\phi\sigma_\phi\rangle$  when  $\phi$  is sampled from a uniform distribution.

## 9. OPTICAL PUMPING INTO A MAXIMALLY ENTANGLED STATE




(a) Example measurement of ground-state fidelity vs number of protocol cycles. Following an initial transient, the steady-state is an oscillation between  $|\Psi^-\rangle$  (after even cycles) and  $|\Psi^+\rangle$  (after odd cycles). In this measurement, the protocol was not yet well calibrated, and the steady-state singlet fidelity for even cycles is below 0.9.



(b) Singlet fidelity  $F(|\Psi^-\rangle)$  versus the number of cycles  $N$  for different initial states. All the initial states converge towards the same final state, with a large overlap with  $\Psi^-$ . However, the input fidelity differences remain resolvable even for  $N = 16$ . The points were connected only to guide the eye. Error bars are typically smaller than the point size.

**Figure 9.16:** Measurements of dynamics of collective optical pumping.



**Figure 9.17:** Singlet fidelity  $F(|\Psi^-\rangle)$  before (green) and after (blue)  $N = 16$  cycles of the protocol. Before the protocol is applied, the fidelity varies sinusoidally with the phase offset of the analysis pulse, up to a maximum of  $F(|\Psi^-\rangle) \approx 0.88$  (the middle point at  $F(|\Psi^-\rangle) > 0.9$  is an outlier from the typical results and could not be reproduced). At  $N = 16$ , the fidelity  $F(|\Psi^-\rangle) \approx 0.93$  is approximately independent for the initial fidelity.

the phase offset in the same range, averaging to  $F(|\Psi^-\rangle) = 0.93$  (blue dashed line) with a standard deviation of 0.01 over all data points. The only data point where we measured reduced output fidelity was the left-most one, which could either indicate that convergence has not yet occurred, or be due to an experimental drift.

In the remainder of this section, we prepare the input states with fidelity  $F(|\Psi^-\rangle) \geq 0.75$  in order to guarantee convergence. While the exact protocol fidelity drifts over time, averaging over this and all the subsequent datasets, we find a steady-state fidelity of

$$F(|\Psi^-\rangle) = 0.93(1),$$


where the error bar is the standard deviation over all the datasets, and represents mainly experimental drifts rather than measurement shot noise.

While the dissipative protocol is significantly slower at preparing entangled states than unitary gates, the measurement described above highlights its potential as a practical *entanglement purification* method. In our case, the fidelity of unitary preparation of  $|\Psi^-\rangle$  is limited to  $\approx 88\%$  due to noise in single-ion addressing. The dissipative preparation of  $|\Psi^-\rangle$  improves the fidelity to  $\approx 93\%$ , but takes over  $40\times$  longer. The best available method of singlet preparation is therefore to combine the two, and append a small number of collective optical pumping cycles at the end of the calibrated unitary operation. Beyond the potential to improve the fidelity, this method has an added advantage of added resilience to certain sources of miscalibrations, such as the phase offset discussed above.

When I first tested the protocol, I became convinced it is very good improving  $|\Psi^-\rangle$  preparation fidelity in the presence of errors in the common-mode detuning  $\delta_{\text{com}}$  of the MS

## 9. OPTICAL PUMPING INTO A MAXIMALLY ENTANGLED STATE

pulse. The data demonstrating that effect is shown in Fig. 9.18. However, I afterwards



**Figure 9.18:** The measured singlet fidelity  $F(|\Psi^-\rangle)$  in the presence of an artificial common-mode detuning error in the MS interaction. The fidelity at the end of the collective optical pumping protocol (blue) exhibits significantly less error sensitivity than the fidelity after a unitary preparation (red). However, it appears that the unitary preparation suffered from an avoidable calibration error (see main text).

discovered that the strong dependence of the unitary preparation fidelity on  $\delta_{\text{com}}$  was not expected in the simulations. As it turned out, the infidelity was predominantly caused by a self-inflicted phase tracking error. Recall that, in order to prepare  $|\Psi^-\rangle$ , we first perform an MS gate followed immediately by a  $\pi/2$  “analysis” pulse (Sec. 9.3.4). The state preparation error is influenced by the difference between the analysis pulse phase  $\phi$  and the phase of the Bell-state after an MS gate. However, part of this difference is caused by our phase tracking method; while the phase of the analysis pulse is tracked in the frame rotating at the bare qubit frequency  $\omega_0$ , resulting in  $\phi \propto \omega_0 t$ , the spin phase of the MS drive (which determines the phase of the Bell state in absence of other errors) is calculated in the frame which includes the common-mode frequency offset, and thus  $\phi_s \propto (\omega_0 + \delta_{\text{com}})t$  (Eq. (9.16)). As a result, deliberate changes in  $\delta_{\text{com}}$  result in a phase offset of

$$\Delta\phi = \phi - \phi_s = \delta_{\text{com}}(t_{\text{MS}} + \tau),$$

where  $t_{\text{MS}}$  is the MS pulse length and  $\tau$  is the delay between the time when the phase counters are set to zero and the beginning of the state initialisation sequence. When taking the data, not only did I forget to correct for this phase offset, I did not even explicitly keep track of the decay  $\tau$  between the phase reset and the MS drive. Indeed, the data in Fig. 9.18 implies a phase error of  $\Delta\phi \approx 90$  deg for  $\delta \approx 0.8$  kHz and  $t_{\text{MS}} = 75$   $\mu\text{s}$ , suggesting  $\tau \approx 240$   $\mu\text{s}$ .

As a result, we both can and cannot claim that the protocol is beneficial in the presence of common-mode detuning errors. On one hand, the data is insufficient to determine whether the fidelity achieved with the dissipative sequence exceeds the fidelity of the unitary


preparation when the latter is done without avoidable errors. On the other hand, the benefit of the dissipative sequence is that there are no phase errors that have to be avoided in the first place!

#### 9.4.4. Understanding errors

As discussed in Sec.9.2.2, off-resonant excitations and residual spin-motion entanglement both contribute to the protocol error. While their individual effects are difficult to distinguish, we capture their combined influence through the bit-flip probability  $p$ . In the correlated bit-flip model, the steady-state error  $\epsilon \approx 3.2p$ . The bit-flip probability  $p$  can be measured as discussed in Sec. 9.3.2.

We expect  $p$  to increase for larger motional mode temperatures (Sec. 2.5.1). This is the case for both the MS drive errors - which increase with gate mode temperature and spectator mode temperature - and for the spectator mode excitation. As a result, the protocol exhibits no true steady-state for finite heating rates, and the “instantaneous steady-state fidelity” always decreases over time.

In order to quantify this effect, we measure  $p$  following a gap time between laser cooling and state initialisation. In the gap time, we either apply  $N$  cycles of the dissipative protocol, or insert a wait time of the same length. In the measurement, we typically assume  $p = p(\downarrow)$  as defined in Sec. 9.3.2, i.e. that the mode frequencies are such that the excitation probability of  $|\uparrow\rangle$  by the drive (A) is negligible.



**Figure 9.19:** The measured bit-flip error  $p$  after  $N$  cycles of the protocol. Different colors display datasets collected at different periods of the day, and dashed lines indicate their individual linear fits. All the datasets point to the same trend of  $p \approx 0.01$  at  $N = 0$  increasing to  $p \approx 0.02$  for  $N = 16$  cycles

Fig. 9.19 shows three such measurements collected during the data acquisition period, together with linear fits. For each dataset, we find that  $p \approx 0.01$  for a ground-state cooled system increases to  $p \approx 0.02$  for  $N = 16$  cycles. The measured steady-state protocol error  $\epsilon \approx 0.07(1)$  is consistent with the prediction of  $\epsilon \approx 3.2 \times p$ . This implies that, while we

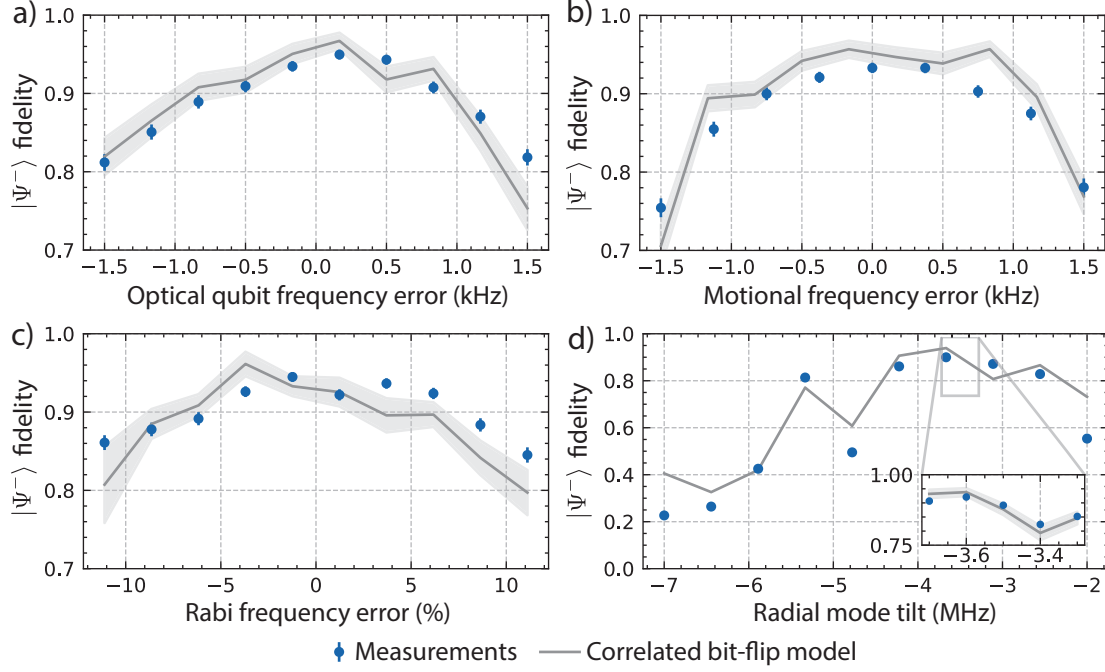
cannot precisely explain what limits the drive (A), the observed bit-flip error is sufficient to explain the protocol error.

Furthermore, we record no difference between  $p$  measured after  $N$  protocol cycles or after a wait-time of the same length. This indicates that the errors are dominated by the background heating rates and not by protocol-induced heating. In principle, the protocol itself can increase the gate mode temperature as follows. Suppose that after the application of the drive (A) the spin and motion are entangled ( $\alpha \neq 0$ ). The subsequent repump step (B) resets the spin, thus changing motional state. Formally, for  $|\alpha| \ll 1$ , tracing out the spin corresponds to applying an oscillator creation or annihilation operator with probability  $p \propto |\alpha|^2$ , which is analogous to a jump operator used to model motional heating (Tab. 2.3). As a result, the next cycle starts with a hotter motional state. Furthermore, a hotter initial state generally leads to larger residual spin-motion entanglement at the end of the drive (A), and thus this effective heating rate increases over time. However, simulations indicate this to be a negligible effect when the drive (A) is phase-modulated, and we indeed record no evidence of it in the experiment.

In order to verify the reach of the correlated bit-flip error model, we deliberately introduced calibration errors into the protocol. For every parameter setting, we measured  $p$  after  $N = 16$  cycles of the protocol, and used it to compute the expected steady-state fidelity. Simultaneously, we measured the steady-state fidelity as described above. Fig. 9.20 shows the comparison between the measured and expected fidelities.

The measured fidelities (blue points) are indeed in very good qualitative agreement with the correlated bit-flip model predictions (grey lines) for a wide range of experimental settings. While systematic deviations remain, our data indicates that this single parameter captures the essential physics of the protocol imperfections. The bit-flip errors and their uncertainties were estimated by averaging multiple individual measurements conducted throughout the day. These generally gave consistent results, with the exception of the data presented in Fig. 9.20 c), where the Rabi frequency  $\Omega$  of the drive (A) was varied. In those experiments, the average power of the  $\lambda = 729$  nm laser changed as well, causing significant frequency fluctuations and erratic behaviour. Fig. 9.20 d) illustrates the effect of the bit-flip error associated with varying the radial mode tilt curvature already discussed in Sec. 9.3.2. Indeed, only a narrow range of radial mode orientations and frequency allows for high-fidelity singlet state preparation.

While the measured errors are in agreement with the correlated bit-flip model, there are two significant qualitative disagreements with the error simulations discussed in Sec. 9.2.3. First, while the ideal protocol exhibits little-to-no sensitivity to the Rabi frequency error, we find the bit-flip probability increases when the Rabi frequency  $\Omega$  is misset. Even in the absence of trap charging, the experimental observations can be explained by the fact that changing  $\Omega$  changes the AC Stark shift of the drive (A), in turn affecting the bit-flip error. This effect is not captured in a two-level simulation, but is present in the experiment due to coupling to dipole transitions. In order for the protocol to achieve the nominal



**Figure 9.20:** Experimentally measured values of  $F(|\Psi^-\rangle)$  (blue dots) compared to the prediction ( $1 - 3.2p$ ) of a correlated bit-flip error model (gray lines). Error bars show a to  $\pm 1\sigma$  confidence interval. In d) the measured values of  $p$  are large, and we can no longer apply the linear approximation. Instead, the steady-state error is evaluated numerically for the measured mean value of  $p(\downarrow)$ ,  $p(\uparrow)$  and the area error of the drive (A). The error model predictions were connected with straight lines to guide the eye.

insensitivity to errors in  $\Omega$ , the AC Stark shift needs to be compensated in a different way, for example with an additional tone of the qubit laser [Häf03]. Second, the fidelity loss measured for non-zero optical qubit or motional frequency errors (Fig. 9.20 a and b) is larger than expected in the simulation. The reason is the gap time between the two pulses of the phase-modulated drive (A) (Sec. 9.3.1). Artificially varying the detunings of the MS Hamiltonian affects the relative phases of the pulses, increasing the errors. Indeed, the experimental results could be reproduced in simulations by adjusting the relative phases of the individual pulses of the drive (A). This error could be eliminated by smarter phase calibration or by smooth phase modulation.

While the protocol fidelity is currently limited to  $F(|\Psi^-\rangle) = 0.93(1)$ , its limitations are setup-specific and do not reflect a fundamental limitation of the method. Coupling to spectator modes could be suppressed by increasing the magnetic field or narrowing the laser spectrum (Sec. 4.4.8). Reducing the heating rates to levels more typical for cryogenic traps with similar ion-electrode distance (Sec. 5.6.4), combined with better shielding of nearby dielectrics, would suppress drifts in  $p$  within each collective pumping sequence and hence allow the protocol to reach its true steady state. Alternatively, mode temperature could be

## 9. OPTICAL PUMPING INTO A MAXIMALLY ENTANGLED STATE

stabilized throughout the protocol by sympathetic cooling [Kie00b]. Therefore, I hope that the protocol can be a practically beneficial for high-fidelity entangled state preparation, especially as a method of purifying lower-fidelity Bell states.



## 10. Summary and Outlook

This thesis described the design and operation of a cryogenic SET experiment. Building upon an existing setup, we upgraded the vacuum system to allow for more modularity and improved performance. The laser system also underwent significant changes, In particular, we created a new narrow-linewidth  $\lambda = 729$  nm beamline for qubit laser delivery via trap-integrated waveguides. This allowed us to generate two-qubit maximally-entangled states with a fidelity of 99.45(13)%, with well-characterised error sources.

We also explored several applications of extending the Hilbert space beyond two levels. Making use of multiple long-lived atomic states, we demonstrated a basic toolbox of single-qutrit control and applied it to violate two NC inequalities. We also demonstrated elements of quantum control of a hybrid Zeeman-optical qubit. In particular, we implemented a novel dissipative entanglement generation scheme, whereby Zeeman qubit entanglement is created by a combination of global optical and radio-frequency fields. This allowed us to create singlet Bell states with 93(1)% fidelity in the steady-state, regardless of the initial state.

### 10.1. Future plans and improvements

The work presented in this thesis opens the door to multiple research avenues. Here, I highlight some ideas about possible improvements to the system, as well as the next-generation traps and experiments.

#### 10.1.1. Experimental setup

The cryogenic setup, both due to its age and modifications, is showing signs of degradation. Replacing it with a new unit that was designed for low vibrations and integration in atomic physics experiments could help with improved performance and more robust assembly.

The ion lifetimes recorded in the setup fall short of the state-of-the-art, and are in fact comparable to many experiments with room-temperature SETs (Sec. 8.1.1). Future experiments should verify whether this is due to high residual vacuum pressure, for example by measuring background-gas collision rates [Aik20; Cla21]. If so, it might be possible to improve the vacuum by only using metal seals, instead of a combination of CF and KF flanges. Alternatively, additional measures might be necessary to reduce outgassing or improve cryopumping of hydrogen, whose peak pressure can occur at a temperature of around 20 K [Bag20]. Finally, any new leaks between the vacuum system and the helium dewar are of great concern, since helium is the least effectively cryopumped common gas.

## 10. SUMMARY AND OUTLOOK

The 4K chamber was electron-beam welded out of OFE copper panels (Sec 3.2). This method offers an easy way of creating low-conductivity joints and hence high B-field shielding factors. However, the shielding should be verified, for example by injecting external B-field noise along different directions [Bra16a]. Depending on the results, future experiments might need to take more care in sourcing high-purity copper.

An effusive oven located at the room-temperature stage has proven to be a more robust source of neutral atoms than either a cryogenic oven or an ablation target (Sec. 3.1.5). As a result, during the writing of this thesis we removed the cryogenic oven altogether. The current downside of the room-temperature oven is that it needs to be removed and re-installed whenever the chamber is opened. This could be helped by using a smaller unit with a fixed oven position. Additional possible improvements include oven temperature control [Bal17] and a shutter [Bra16b] to block unnecessary atom flux into the 4K chamber. This should simultaneously improve the loading rates and alleviate some worries about trap contamination. We also started investigating the possibility of loading from a pre-cooled atomic source [Bru16].

Light delivery through integrated waveguides led to an improvement in optical qubit coherence (Sec. 5.4.3). This is because, while cryogenic vibrations remain significant, the waveguides vibrate in phase with the ions. However, vibrations still appear to be the limiting factor of the Zeeman qubit coherence (Sec. 5.4.4). There, differential oscillations between the ion trap and the permanent magnets placed on the outer vacuum chamber lead to B-field fluctuations. Future experiments can improve this by attaching the magnets to the 4K chamber.

Future experiments could also benefit from simplified signal delivery, especially when increasing the number of connections. High-density wiring based on flexible PCB technology appears to be an attractive option for custom connections [Con16]. Shielding the DC cables from electromagnetic pickup inside the chamber might lead to reductions in technical noise. If the cryostat steel bellows are replaced with rubber ones, it might be possible to decouple the cryostat body from the building earth. This would allow for more grounding flexibility and ground-loop monitoring, potentially reducing the noise as well. However, those improvements might only be beneficial once the trap heating rates are reduced, as current evidence suggests that technical noise is not the limiting factor (Sec. 5.6.4). The anomalously high heating rates observed in our traps with integrated optics have to be addressed through experiments with the fabrication methods. It would be interesting to find out whether the electric-field noise can be suppressed by increasing the trap ground-plane conductivity, or by surface treatment [Hit13; Hit21; Ber21].

### 10.1.2. Optical setup

The dipole laser system, while reliable, could benefit from improved frequency stabilisation. In particular, ion loss is frequently associated with slow drifts of the length of the low-finesse cavity to which the  $\lambda = 397\text{ nm}$  laser is locked, with frequency changes of as much as

10 MHz in 24 hours (Sec. 4.3.3). Unless those drifts can be reduced, we might benefit from a robust and automated cavity length calibration routine. This is particularly important as multiple experiments rely on the same lasers, with frequent disagreements about the optimal settings. It might be beneficial to switch from using variable-length to fixed-length cavities, which offer superior stability.

The spectral noise of the quadrupole laser is one of the dominant error sources for entangling gates (Sec. 8.3.2). Despite significant achievements in spectral filtering of the servo bumps, and injection-locking multiple laser sources, the locked laser linewidth and noise remain to be fully investigated and optimised. The reason this did not happen so far is that the  $\lambda = 729$  nm spectrum was “good enough” for almost all users, and such tests cannot be performed without disrupting multiple experiments. Nonetheless, as other errors are eliminated, decreasing the laser linewidth becomes increasingly urgent. I hope that the measurements shown in this thesis can guide the next users through this process. In order to verify whether injection locking is a scalable method of amplifying the quadrupole laser power, it is imperative to systematically compare the spectral properties of the parent and child lasers (Sec. 4.4.8).

Robust and scalable laser delivery will require increased use of fibre-integrated components instead of free-space routing. To that end, we recently designed test units of fully fiberized<sup>1</sup> injection locking, and in-fibre FNC. These designs, laid out in App. 11.5, make use of integrated photodetectors, laser diodes, AOMs and circulators to create integrated systems that require no manual alignment and are robust to external noise. Preliminary tests of the FNC unit at  $\lambda = 729$  nm found a power transfer efficiency of  $\approx 30\%$  (corrected for the input fibre coupling loss). For fibre input powers as low as 200  $\mu$ W, the lock beat note was visible with signal-to-noise ratio (SNR) of  $> 30$  dB. The design is now being tested by A. Mestre and M. Marinelli for locking performance. The first tests of the in-fibre injection setup, performed recently by A. Vasquez, demonstrated that stable in-fibre injection locking at  $\lambda = 729$  nm was possible with seed power above approximately 10  $\mu$ W. While the use of fibre components decreases the available output power from  $\approx 40$  mW to  $\approx 10$  mW, the benefit of rapid assembly combined with (still unconfirmed) noise resilience may nonetheless make it worthwhile. Furthermore, optical losses can be reduced by splicing fibre components instead of relying on fibre connectors.

### 10.1.3. Hybrid qubit operation

In addition to the collective optical pumping into a maximally entangled state described in this thesis, preliminary experiments described in the thesis of C. Zhang [Zha21] demonstrate how the MS interaction can be combined with shelving pulses to perform two-qubit gates on the Zeeman qubit. Future experiments should extend this toolbox to demonstrate practical advantages of this “hybrid qubit” encoding and manipulation paradigm. One possible research direction is to use a decoherence-free-subspace qubit encoding, where the logical

---

<sup>1</sup> Even though this word is made up, I wonder what the British spelling would be. Fibreised?

## 10. SUMMARY AND OUTLOOK

states are encoded into maximally entangled states of two Zeeman qubits:

$$\begin{aligned} |0\rangle_L &= (|\downarrow\uparrow\rangle + |\uparrow\downarrow\rangle)/\sqrt{2} \\ |1\rangle_L &= (|\downarrow\uparrow\rangle - |\uparrow\downarrow\rangle)/\sqrt{2}. \end{aligned}$$

This would create a very robust quantum memory, beneficial for example for transport sequences in the QCCD architecture. This encoding is also robust to global errors, and in fact, single-ion addressing is necessary to perform single-qubit rotations between  $|0\rangle_L$  and  $|1\rangle_L$  [Aol07; Mon09]. While we demonstrated such single-qubit gates indirectly by exploiting the differential AC Stark shift in Sec. 9.3.5, that method is noisy and non-scalable. Thus, we recently installed fast waveform generators that allow the trap potential to be flexibly modified within one experimental sequence. We hope that this will allow us to implement robust and low-noise single-ion addressing by translating the ions across the  $\lambda = 729$  nm laser beam. In addition, improved trap waveform control will hopefully allow for reliable loading of multi-ion chains, following the method in Sec. 8.1.1.

### 10.1.4. Quantum computing with qudits

Using multi-dimensional quantum systems instead of qubits might decrease the number of ions needed to implement certain quantum algorithms [Gok19]. Multi-level systems also allow for an implementation of the Toffoli gate with fewer operations [Lan09] and offer resource advantages for quantum error correction [Cam14]. However, while single-qudit operations are relatively straightforward to implement with a single ion, multi-qudit entangling gates have not yet been implemented in a trapped-ion system<sup>2</sup>. As proven in [Low19], a two-qudit entangling gate for a pair of  $d$ -level systems can be created with a Hamiltonian that generalises the MS interaction. Experimentally, this can be achieved by applying  $2 \times d$  frequencies onto the laser field, each detuned from the motional sideband of a different carrier transition. However, this implementation may eliminate certain advantages of the enlarged Hilbert space, since the gate duration and error rates will both grow with increasing dimension  $d$ . Furthermore, the qudit MS interaction does not by itself prepare maximally entangled states [Low19], putting a question mark over the efficiency of the universal gate set it is part of. Further work is therefore needed to establish whether there are practical algorithmic advantages to multi-level operations with trapped ions. It would be also extremely valuable for a trapped-ion experiment to perform and benchmark a two-qudit entangling gate. This would help understand the true overheads and challenges of multi-level control. Furthermore, it would clarify whether simultaneous multi-level interactions are beneficial for qudit-based quantum logic, or whether it is better to implement multi-qudit gates simply as sequences pulses addressing one transition at a time.

---

<sup>2</sup> C. Senko *et al* demonstrated trapped-ion two-qutrit entanglement in [Sen15]. However, the Hamiltonian they employed is only valid within a restricted subspace, and therefore cannot be used as part of a universal gate set.

### 10.1.5. Contextuality tests

The contextuality tests described in this thesis suffer from several drawbacks, notably the lack of a clear way to address the compatibility loophole (Sec. 6.5). Future theoretical work is needed to provide a clear interpretation of experimental contextuality tests. Nonetheless, given that the Bell tests have firmly established that the universe is not governed by a local deterministic theory, further contextuality tests lack a strong theoretical motivation (Sec. 6.6.4). As discussed in Sec. 6.6.5, local contextuality tests cannot be used to certify the security of a random number generator, and thus there is also no strong technological motivation behind them either. My current assessment is that local contextuality tests should not be performed until somebody proposes a plausible reason to do them. The only exception is perhaps experimental tests performed within the Spekkens’s operation framework of contextuality (Sec. 6.5.2), which is designed to experimentally address the question of whether the known formalism of quantum mechanics is the ultimate description of the correlations obtained in non-classical scenarios [Maz17].

### 10.1.6. Integrated optics

Following the successful demonstration of single and two-ion optical qubit operations through integrated waveguides, a number of interesting research avenues await implementation. While the work discussed in this thesis only delivered the qubit laser at  $\lambda = 729$  nm through the waveguides, it would be extremely beneficial to integrate all the lasers necessary for trapping and qubit control. This is already possible in our experiment for the repump lasers at  $\lambda = 854$  nm and  $\lambda = 866$  nm, and the traps discussed in this thesis already have dedicated grating couplers for those wavelengths. However, in  $\text{Si}_3\text{N}_4$ , large absorption at low wavelengths prohibited us from integrating waveguides for cooling and readout at  $\lambda = 397$  nm, as well as the photoionization light. While [Nif20] adapted the dielectric deposition process and demonstrated light delivery through  $\text{Si}_3\text{N}_4$  waveguides at wavelengths as short as  $\lambda = 405$  nm, this occurred at a significant loss of 10 dB/cm. An alternative solution is to use a different waveguide material to deliver light at blue and UV wavelengths. G. West *et al* have shown in [Wes19] a fabrication process for alumina waveguides, with propagation losses below 1.5 dB/cm at  $\lambda = 405$  nm and 3 dB/cm at  $\lambda = 370$  nm. Thus, low-loss delivery of all the wavelengths necessary for  $^{40}\text{Ca}^+$  can be achieved by combining alumina and  $\text{Si}_3\text{N}_4$  waveguides on a single chip.

The ability to integrate UV-compatible waveguides into the trap chip would allow for fluorescence collection through the waveguides as well. Unavoidable optical losses necessarily imply that this method would reduce the photon flux compared to integrating a photon detector into the trap directly underneath the ion. However, readout through trap-integrated waveguides offers several advantages. High numerical aperture can be achieved by trapping an ion directly above the waveguide and bringing it close to the surface. By comparison, recent attempts at operating an SET with an integrated single-photon avalanche photodetector revealed very large ion heating in the vicinity of the detector

## 10. SUMMARY AND OUTLOOK

[Set21]. Integrated waveguides can be used to deliver photons to either an on-chip or an off-chip detector. The increased distance between the ion and the detector could allow the use of SNSPDs without the performance loss associated with the trap RF or the ion excitation caused by the detector pulses [Tod20]. The readout light at  $\lambda = 397\text{ nm}$  could be collected through a separate grating coupler, or with the same beam path that is to deliver the light. In the latter case, the photons travelling in different directions could be separated (with additional loss) by cross-polarization or using a circulator.

It remains to be demonstrated that SETs with integrated light delivery can operate at cryogenic temperatures with low drifts and heating rates. We found this to be a major issue, with significant light-induced charging, and electric-field noise strength one to two orders of magnitude worse than the state-of-the-art. Encouragingly, Ivory *et al* [Ivo20] recently characterised a room-temperature SET with integrated waveguides, finding that the electric-field noise was on par with other room-temperature traps, and did not measurably increase in the vicinity of the grating coupler. This was particularly encouraging as the ion-electrode distance was only  $20\text{ }\mu\text{m}$ , the lowest of any published experiment that I am aware of. Next experiments should verify whether the electric-field noise observed at room temperature decreases substantially in cryogenic environments.

Light-induced charging is an outstanding issue also in [Ivo20]. The most obvious method of mitigating this noise is to deposit an ITO coating over the exposed dielectrics. This was the approach used in [Nif20], which reported no observable charging from trap-integrated light. However, we only recorded significant charging with milliwatt-level power emerging out of the grating coupler, significantly exceeding the powers available in [Nif20] due to their large optical losses. Thus, the degree of stray field suppression of metal coatings remains to be verified. Future experiments should aim to find a thickness of the ITO layer that reduces frequency fluctuations to negligible levels while minimising optical losses. ITO was already demonstrated as a viable trap material in [Elt12], with 60% transmission measured for a quartz block with a  $400\text{ nm}$ -thick ITO coating. Alternatively, a thin semi-transparent layer (a few nm) of a high-conductivity metal such as gold can be used instead of ITO [Sto03].

All the chip-integrated light delivery demonstrations to date involved on-chip routing and off-chip modulation. In order to scale the approach to large multi-zone registers, it is necessary to integrate active modulation onto the chip. This is challenging as the waveguide material must allow for modulation in addition to the already discussed requirement of low loss for a broad range of UV-to-IR wavelengths. Furthermore, trapped-ion experiments require a modulator with high extinction ratio (Sec. 5.4.1), switching speeds of  $\sim 1\text{ MHz}$ , and cryogenic operation places additional constraints on power consumption. Many of the promising methods of on-chip intensity modulation are based on the Mach-Zehnder interferometer configuration [Ree10], where the refractive index of one of the arms is modulated, affecting the relative phase of the interfering paths, and in turn the relative intensities in the output ports. In lithium niobate, it is possible to change the refractive index at high speed by electro-optic modulation. Recent work by B. Desiatov *et al* [Des19]

demonstrated an intensity EOM based on thin-film lithium niobate with loss-loss operation at visible wavelengths (here at  $\lambda = 636$  nm) and an extinction ratio of  $\approx 30$  dB. While  $\text{Si}_3\text{N}_4$  does not allow for electro-optic modulation, the refractive index can be adjusted by changing the local temperature or the waveguide strain. The former thermo-optic modulation of the refractive index is likely insufficient for trapped-ion applications due to limited switching rates ( $\approx 1$  kHz) and large power consumption ( $\approx 40$  mW per modulator in  $\text{Si}_3\text{N}_4$ ) [Arr21]. On the other hand, recent work from Dong *et al* [Don21] demonstrated high-speed  $\text{Si}_3\text{N}_4$  strain modulation using aluminium nitride piezoelectric actuators. This allowed them to create an intensity modulator with up to 100 MHz switching rate and an extinction ratio of 30 dB. At a switching rate of 1 MHz, the recorded power consumption was below 200  $\mu\text{W}$ , with only 6 nW necessary to hold the state. However, the interferometer had a relatively large footprint of  $\approx 1$  mm<sup>2</sup>. An alternative approach to intensity modulation uses micromechanical (MEMS) actuators. Recently S. Gyger *et al* [Gyg21] demonstrated a MEMS coupler with over 100 kHz switching rate, 28 dB extinction ratio and power consumption below 75  $\mu\text{W}$ . While slower than the electro-optic and strain-based modulators discussed above, this approach may be sufficient for trapped-ion applications, and has the advantage of a small footprint ( $\approx 20 \mu\text{m} \times 80 \mu\text{m}$ ).

I am really looking forward to the first demonstration of ion trapping with fully integrated light delivery and readout. Such a device has the potential to completely change how ion traps are designed and tested. When optical access is not required, ions can be confined *inside* fully enclosed structures, for example “sandwiched” between two planar chips [Rag20]. This can improve trap depths and provide shielding from external noise. Furthermore, with fully integrated light delivery, it is possible to control ions trapped very close to the chip surface without increasing laser scatter and cross-talk associated with free-space addressing. Thus, the ion-electrode distance  $d$  can be reduced, with the minimum value set by the maximum tolerable heating rate [Ste04]. This can be very advantageous for implementing multi-ion microwave gates, where the sideband coupling strength achievable for a given current scales as  $d^{-2}$  [Osp11]. Full integration also allows multiple ion traps to be operated and tested in parallel. Thus, it might be possible to speed up the ion trap development cycle by automating the testing process.

### 10.1.7. Dissipation engineering

The demonstration of collective optical pumping into a maximally entangled state described in Chapter 9 is the first example of a dissipative process generating entangled states with  $> 90\%$  fidelity. I hope that due to its simplicity and noise resilience, the protocol can be practically useful, and that next-generation experiments can implement it to obtain  $> 99\%$  Bell-state fidelities. This should be reasonably straightforward in a setup with reduced spectral crowding and lower heating rates. While the analysis presented in this thesis was focused on the trapped-ion implementation, the protocol is general and we anticipate it might be applied in a wide range of platforms where suitable couplings are available. For

## 10. SUMMARY AND OUTLOOK

example, in nitrogen-vacancy centres, the action of the MS gate may be replaced by direct spin-spin interactions [Maz11]. In neutral atom platforms, Rydberg anti-blockade provides a possible mechanism to implement the collective excitation [Mit20]. In superconductors, suitable two-body interactions may be realised by parametric driving [Dou20].

Future dissipation engineering experiments should also aim to directly create entangled states of more than two qubits. For example, D. Cole *et al* recently proposed in [Col21] a method to create three-qubit “W states” [Häf05] based on symmetry engineering. In the proposal, the target state can be created without the need for single-ion addressing, and at competitive timescales. I hope that this proposal can be implemented in the nearby future, and extended to larger qubit registers.

Finally, while the existing dissipative entanglement protocols can be beneficial in the presence of certain errors, they have not been in practice. As a result, the fidelities of entangled states prepared through dissipative methods have not been competitive with those achieved using unitary control. The big outstanding challenge of dissipative engineering is thus to achieve a “dissipative advantage” by demonstrating the preparation of a non-trivial state with errors lower than other approaches. Alternatively, the benefit could be achieved by demonstrating high-quality entanglement generation with high tolerance to relevant sources of drifts [Bal20].



# 11. Appendix

## 11.1. Clebsch-Gordan coefficients in $^{40}\text{Ca}^+$

The Clebsch-Gordan coefficients describe the relative branching ratios of different sub-levels of the same fine-structure state.

	$m_j$	$S_{1/2}$		$P_{1/2}$		$P_{3/2}$			
		$-1/2$	$+1/2$	$-1/2$	$+1/2$	$-3/2$	$-1/2$	$+1/2$	$+3/2$
$S_{1/2}$	$-1/2$			$1/3$	$2/3$	$1$	$2/3$	$1/3$	
	$+1/2$			$2/3$	$1/3$		$1/3$	$2/3$	$1$
$D_{3/2}$	$-3/2$	$1/5$	$4/5$	$1/2$		$3/5$	$2/5$		
	$-1/2$	$2/5$	$3/5$	$1/3$	$1/6$	$2/5$	$1/15$	$8/15$	
	$+1/2$	$3/5$	$2/5$	$1/6$	$1/3$		$8/15$	$1/15$	$2/5$
	$+3/2$	$4/5$	$1/5$		$1/2$			$2/5$	$3/5$
$D_{5/2}$	$-5/2$	$1$				$2/3$			
	$-3/2$	$4/5$	$1/5$			$4/15$	$2/5$		
	$-1/2$	$3/5$	$2/5$			$1/15$	$2/5$	$1/5$	
	$+1/2$	$2/5$	$4/5$				$1/5$	$2/5$	$1/15$
	$+3/2$	$1/5$	$4/5$					$2/5$	$4/15$
	$+5/2$		$1$						$2/3$

## 11.2. Injection lock alignment

An injection lock is set up as follows. A child diode is placed in a temperature-controlled mount and turned on. Its output is collimated by an aspherical lens to roughly match the seed beam size. An optical isolator<sup>1</sup> is placed at the centre of the beam, with the horizontal cube facing away from the diode. A HWP between the diode and the isolator is adjusted to maximise transmission through the isolator. Then, the seed parent beam can be aligned. First, its angle should be roughly aligned until a clean beam reaches the diode. Then, its polarization can be rotated to maximise the amount of power reaching the diode. The next, finer adjustment, can be done by turning up the diode current above the lasing threshold. Due to polarization impurity, a small fraction of the child diode output will exit the isolator at the seed port. The seed beam should then be overlapped with this mode. By this point, it should be possible to observe and fine-tune diode injection. The most straightforward

<sup>1</sup> The isolator should be already aligned to maximise isolation at the desired wavelength

signature of that is the amplification of the diode output power close to the lasing threshold (typically by many orders of magnitude). The beam overlap can be therefore fine-tuned by maximising the diode output power close to its lasing threshold. Once the modes are overlapped, the diode can be injected also above the lasing threshold whenever the seed wavelength is resonant with the diode cavity. For any diode temperature, this occurs at a series of discrete and approximately equally diode currents. We can find these points by monitoring the output light on a scanning Fabry-Perot cavity. For AR-coated diodes, it is also possible to optimise the injection by monitoring the output power alone.

### 11.3. Undesired FNC back-reflections

There are a few common FNC failure modes that can creep up on a unlucky student. The most important one is a back-reflection from the (angled) fibre input. Since it is reflected at an angle, it is easy to spot and aperture if necessary. The tricky ones originate from objects after the retro-reflector. I have seen this happen in two circumstances:

1. The fiber output was used to injection-lock a laser diode. A fraction of the diode's output was leaking back into the fiber, and since it was aligned very well, the light was well-coupled into the fibre. When the diode was turned off, or turned on but not injected well, this only produced incoherent light and did not disturb the injection. However, whenever the diode was injected, this leakage would overwhelm the correct retro-reflection signal and de-stabilise the FNC lock
2. The fiber output was used in another FNC setup. From Fig. 4.11, one can easily see that any FNC setup will reflect a significant fraction of the input light back at the source, some at the incident frequency, and some shifted by  $2f_{\text{AOM}}$ . The effect on daisy-chained FNCs could be two-fold. The light reflected at the incident frequency will act as a second retro-reflector, and could thus de-stabilise the first lock as described above. The light shifted in frequency will create a second beatnote at the first photodiode, which will overlap in frequency with the first beat-note unless the AOMs operate at different frequencies.

In the end, the safest choice is always to put an optical isolator after the retro-reflector. Otherwise, one should measure the power and frequency of the additional reflections, and align the optical elements to ensure there is one and only one strong retro-reflection present. Furthermore, it is a good idea to run different FNCs at different AOM frequencies.

### 11.4. AOM extinction ratio

Fundamentally, the extinction ratio of an AOM is limited by the mode overlap between the 0th and 1st order beam at the shutter. For Gaussian beams and aberration-free optics,

this can be a minor correction<sup>2</sup>. In practice, the extinction ratio is limited by undesired reflections and scattering, as well as unavoidable beam profile distortions. Ideally, leakage should be measured and corrected if necessary. However, the following design principles minimise the risk of having a problem in the first place:

- Use an AOM with high centre frequency  $f_{\text{RF}}$  and a lens with long focal length  $f$  to obtain a large beam separation  $\Delta x$
- Place the shutter close to the mirror, where the beams are smallest
- Carefully align the shutter along  $x$ , blocking as large of a region as possible. If in doubt, centre its edge closer to the diffracted spot.
- Use a black shutter, and angle it with respect to the optical axis to avoid back-reflections
- Place the QWP after (and not before) the lens. Then, any reflections from the lens will not have their polarization rotated, thus not coupling to the output port.

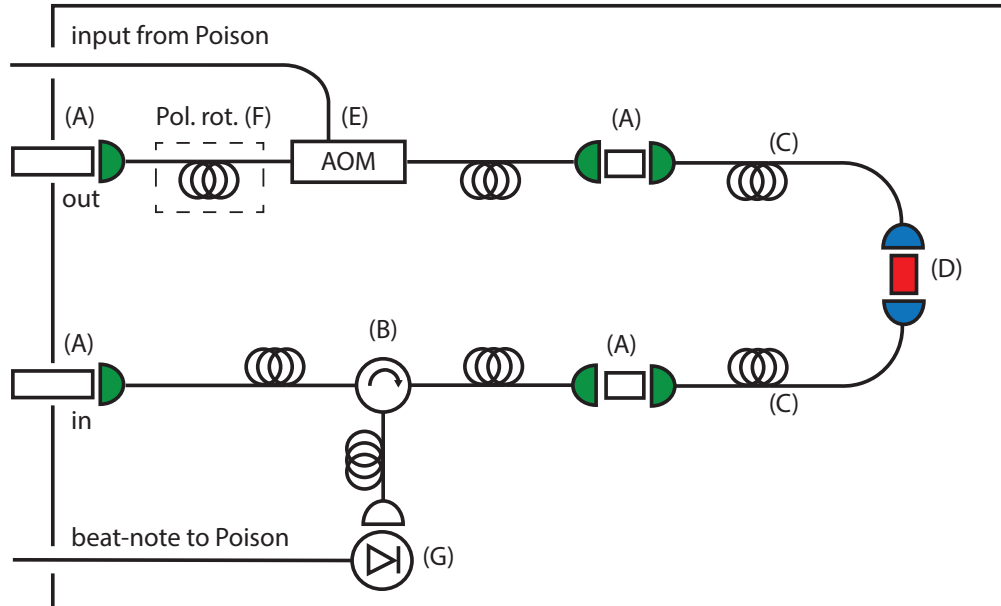
## 11.5. In-fibre modules for $\lambda = 729$ nm laser

Fig. 11.1 presents a schematic layout of the test modules for the in-fibre FNC and in-fibre injection locking for  $\lambda = 729$  nm. Unless otherwise specified, all the optical fibres are SM SM600 pathcords with APC connectors on either end. The remaining components are listed in Tab. 11.1, which serves as a legend for Fig. 11.1.

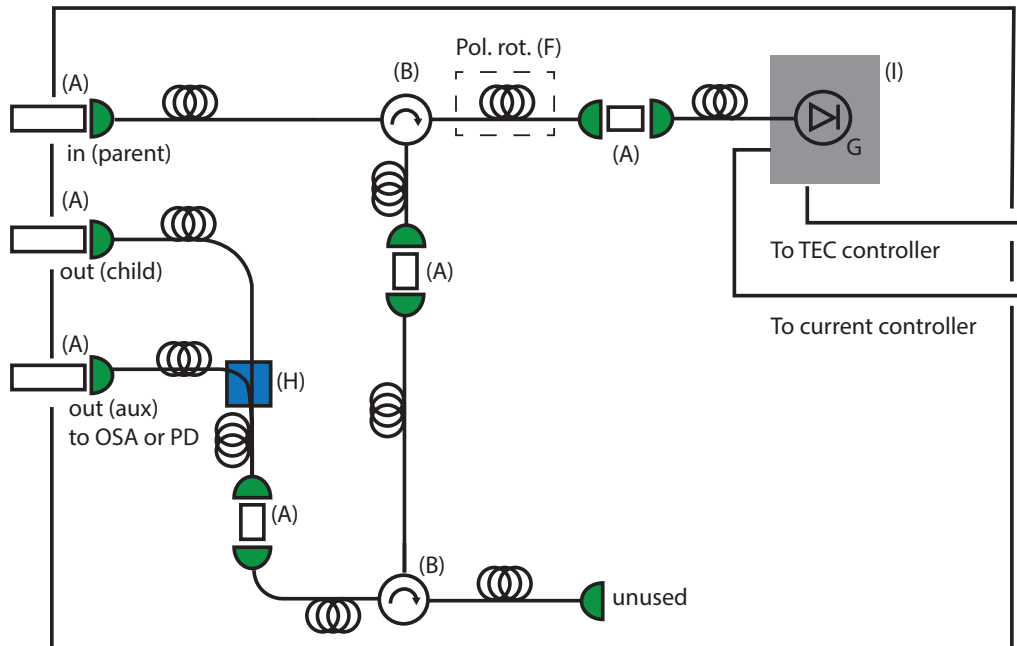
---

<sup>2</sup> For example, a  $\lambda = 854$  nm gaussian beam of  $\phi = 1$  mm diameter, pulsed by a double-pass setup with  $f_{\text{RF}} = 2\pi \times 80$  MHz,  $f = 7.5$  cm,  $v_S = 4260$  m/s and shutter-mirror distance  $d = f/2$  has its extinction ratio limited to  $-2000$  dB by the mode overlap.

## In-fibre fibre-noise cancellation



## In-fibre diode injection lock



**Figure 11.1:** Component layout for the in-fibre FNC unit (top) and in-fibre diode injection lock (bottom). Tab. 11.1 gives a legend for the components described by letters (A) - (J) in the figure. Fibre connectors are coloured following a standard convention, where green and blue denote APC and PC connectors respectively.

**Table 11.1:** List of components used in the test units of in-fibre FNC and in-fibre injection locking.

Letter	Component	Supplier / part number
A	Mating sleeve (APC/APC)	Thorlabs ADAFC3
B	Fibre circulator	Ascennta VCIR-3-729 (custom order)
C	Fibre patchcord APC/PC	Thorlabs P5-630A-PCAPC-1
D	Partially reflective PC/PC connector	Thorlabs ADAF1 (misaligned for $\approx 1 - 3\%$ reflection)
E	Fibre AOM	Gooch&Housego fibre-Q (729 nm, 150 MHz)
F	Fibre polarization controller	Thorlabs CPC900
G	Photodetector	Electro-Optics Technology Inc. ET-2060
H	Fibre splitter (99/1)	Thorlabs TW670R1A1
I	Pigtailed laser diode mount	Thorlabs LDM9LP
J	Pigtailed laser diode	Thorlabs LP730-SA15-SP



# Bibliography

- [Aar10] S. Aaronson and A. Arkhipov. “The Computational Complexity of Linear Optics”. *Proc. Annu. ACM Symp. Theory Comput.* 333–342 (2010). Cited on page 4.
- [Abb10] A. A. Abbott, C. S. Calude, and K. Svozil. “A Quantum Random Number Generator Certified by Value Indefiniteness”. *Math. Struct. Comput. Sci.* **24** (2010). Cited on page 228.
- [Abb12] A. A. Abbott, C. S. Calude, J. Conder, and K. Svozil. “Strong Kochen-Specker theorem and incomputability of quantum randomness”. *Phys. Rev. A* **86**, 62109 (2012). Cited on page 228.
- [Abr11] S. Abramsky and A. Brandenburger. “The Sheaf-Theoretic Structure Of Non-Locality and Contextuality”. *New J. Phys.* **13** (2011). Cited on page 223.
- [Abr17] S. Abramsky, R. S. Barbosa, and S. Mansfield. “The contextual fraction as a measure of contextuality”. *Phys. Rev. Lett.* **119** (2017). Cited on pages 223, 228.
- [Ací06a] A. Acín, S. Massar, and S. Pironio. “Efficient quantum key distribution secure against no-signalling eavesdroppers”. *New J. Phys.* **8**, 126 (2006). Cited on page 201.
- [Ací06b] A. Acín, N. Gisin, and L. Masanes. “From Bell’s theorem to secure quantum key distribution”. *Phys. Rev. Lett.* **97** (2006). Cited on page 201.
- [Ací07] A. Acín, N. Brunner, N. Gisin, S. Massar, S. Pironio, and V. Scarani. “Device-independent security of quantum cryptography against collective attacks”. *Phys. Rev. Lett.* **98** (2007). Cited on page 198.
- [Aha04] D. Aharonov, W. van Dam, J. Kempe, Z. Landau, S. Lloyd, and O. Regev. “Adiabatic Quantum Computation is Equivalent to Standard Quantum Computation”. *SIAM Rev.* **50**, 755–787 (2004). Cited on page 287.
- [Ahr13] J. Ahrens, E. Amsellem, A. Cabello, and M. Bourennane. “Two Fundamental Experimental Tests of Nonclassicality with Qutrits”. *Sci. Rep.* **3**, 2170 (2013). Cited on pages 252, 253.
- [Aik20] Y. Aikyo, G. Vrijsen, T. W. Noel, A. Kato, M. K. Ivory, and J. Kim. “Vacuum Characterization of a Compact Room-temperature Trapped Ion System”. *Appl. Phys. Lett.* **117** (2020). Cited on pages 259, 331.

## BIBLIOGRAPHY

- [All12] D. T. Allcock, T. P. Harty, H. A. Janacek, N. M. Linke, C. J. Ballance, A. M. Steane, D. M. Lucas, R. L. Jarecki, S. D. Habermehl, M. G. Blain, D. Stick, and D. L. Moehring. “Heating rate and electrode charging measurements in a scalable, microfabricated, surface-electrode ion trap”. *Appl. Phys. B Lasers Opt.* **107**, 913–919 (2012). Cited on page 259.
- [Alo14] C. Alonso-Ramos, P. Cheben, A. Ortega-Moñux, J. H. Schmid, D.-X. Xu, and I. Molina-Fernández. “Fiber-chip grating coupler based on interleaved trenches with directionality exceeding 95%”. *Opt. Lett.* **39**, 5351 (2014). Cited on page 138.
- [Ami10] J. M. Amini, H. Uys, J. H. Wesenberg, S. Seidelin, J. Britton, J. J. Bollinger, D. Leibfried, C. Ospelkaus, A. P. Vandevender, and D. J. Wineland. “Toward scalable ion traps for quantum information processing”. *New J. Phys.* **12**, 033031 (2010). Cited on page 12.
- [Ami11] J. M. Amini, J. Britton, D. Leibfried, and D. J. Wineland. “Micro-Fabricated Chip Traps for Ions”. In: *Atom Chips*. Weinheim, Germany: Wiley-VCH Verlag GmbH & Co. KGaA, 2011, pages 395–420. Cited on pages 10, 11.
- [Ams12] E. Amsellem, L. E. Danielsen, A. J. López-Tarrida, J. R. Portillo, M. Bourennane, and A. Cabello. “Experimental Fully Contextual Correlations”. *Phys. Rev. Lett.* **108**, 200405 (2012). Cited on page 239.
- [An19] D. An, C. Matthiesen, E. Urban, and H. Häffner. “Distance scaling and polarization of electric-field noise in a surface ion trap”. *Phys. Rev. A* **100** (2019). Cited on pages 12, 181, 183.
- [Ant09] P. B. Antohi, D. Schuster, G. M. Akselrod, J. Labaziewicz, Y. Ge, Z. Lin, W. S. Bakr, and I. L. Chuang. “Cryogenic ion trapping systems with surface-electrode traps”. *Rev. Sci. Instrum.* **80**, 013103 (2009). Cited on pages 12, 55, 259.
- [Aol07] L. Aolita, L. Davidovich, K. Kim, and H. Häffner. “Universal quantum computation in decoherence-free subspaces with hot trapped-ions”. *Phys. Rev. A - At. Mol. Opt. Phys.* **75** (2007). Cited on page 334.
- [App01] D. M. Appleby. *Nullification of the Nullification*. 2001. arXiv: 0109034 [quant-ph]. Cited on page 221.
- [APS17] APS. *Physical Review Journals - Copyright Policies and FAQ*. 2017. (Visited on 05/18/2021). Cited on page 231.
- [Ara] M. Araujo. *Understanding Bell’s theorem part 2: the nonlocal version / More Quantum*. (Visited on 03/31/2020). Cited on page 197.
- [Ara13] M. Araújo, M. T. Quintino, C. Budroni, M. T. Cunha, and A. Cabello. “All noncontextuality inequalities for the n-cycle scenario”. *Phys. Rev. A - At. Mol. Opt. Phys.* **88**, 022118 (2013). Cited on page 212.



- [Ara16] M. Araujo. *Understanding Bell's theorem part 1: the simple version / More Quantum*. 2016. (Visited on 03/31/2020). Cited on page 195.
- [Ari96] E. Arimondo. "V Coherent Population Trapping in Laser Spectroscopy". *Prog. Opt.* **35**, 257–354 (1996). Cited on page 50.
- [Arr21] J. M. Arrazola, V. Bergholm, K. Brádler, T. R. Bromley, M. J. Collins, I. Dhand, A. Fumagalli, T. Gerrits, A. Goussev, L. G. Helt, J. Hundal, T. Isacsson, R. B. Israel, J. Izaac, S. Jahangiri, R. Janik, N. Killoran, S. P. Kumar, J. Lavoie, A. E. Lita, D. H. Mahler, M. Menotti, B. Morrison, S. W. Nam, L. Neuhaus, H. Y. Qi, N. Quesada, A. Repington, K. K. Sabapathy, M. Schuld, D. Su, J. Swinerton, A. Száva, K. Tan, P. Tan, V. D. Vaidya, Z. Vernon, Z. Zabaneh, and Y. Zhang. "Quantum circuits with many photons on a programmable nanophotonic chip". *Nature* **591**, 54–60 (2021). Cited on page 337.
- [Asa19] W. Asavanant, Y. Shiozawa, S. Yokoyama, B. Charoensombutamon, H. Emura, R. N. Alexander, S. Takeda, J. ichi Yoshikawa, N. C. Menicucci, H. Yonezawa, and A. Furusawa. "Generation of time-domain-multiplexed two-dimensional cluster state". *Science (80-. )*. **366**, 373–376 (2019). Cited on page 3.
- [Bab18] R. Babbush, C. Gidney, D. W. Berry, N. Wiebe, J. McClean, A. Paler, A. Fowler, and H. Neven. "Encoding Electronic Spectra in Quantum Circuits with Linear T Complexity". *Phys. Rev. X* **8** (2018). Cited on page 3.
- [Bab21] R. Babbush, J. R. McClean, M. Newman, C. Gidney, S. Boixo, and H. Neven. "Focus beyond Quadratic Speedups for Error-Corrected Quantum Advantage". *PRX Quantum* **2**, 10103 (2021). Cited on page 3.
- [Bag20] V. Baglin. *Cryopumping and Vacuum Systems*. 2020. arXiv: 2006.01574. Cited on page 331.
- [Bag21] G. Baggio, F. Ticozzi, P. D. Johnson, and L. Viola. *Dissipative Encoding of Quantum Information*. 2021. arXiv: 2102.04531. Cited on page 3.
- [Bal14] C. J. Ballance. "High-Fidelity Quantum Logic in Ca +". PhD thesis. University of Oxford, 2014. Cited on pages 29, 277, 279, 280, 284.
- [Bal16] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas. "High-Fidelity Quantum Logic Gates Using Trapped-Ion Hyperfine Qubits". *Phys. Rev. Lett.* **117**, 060504 (2016). Cited on page 2.
- [Bal17] 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". *Rev. Sci. Instrum.* **89** (2017). Cited on page 332.

## BIBLIOGRAPHY

- [Bal20] H. Ball, M. J. Biercuk, A. Carvalho, J. Chen, M. Hush, L. A. De Castro, L. Li, P. J. Liebermann, H. J. Slatyer, C. Edmunds, V. Frey, C. Hempel, and A. Milne. *Software tools for quantum control: Improving quantum computer performance through noise and error suppression*. 2020. arXiv: 2001.04060. Cited on page 338.
- [Bar04] J. Barrett and A. Kent. “Non-contextuality, finite precision measurement and the Kochen-Specker theorem”. *Stud. Hist. Philos. Sci. Part B - Stud. Hist. Philos. Mod. Phys.* **35**, 151–176 (2004). Cited on pages 221, 222.
- [Bar05] J. Barrett, L. Hardy, and A. Kent. “No signaling and quantum key distribution”. *Phys. Rev. Lett.* **95**, 010503 (2005). Cited on page 201.
- [Bar11] J. T. Barreiro, M. Müller, P. Schindler, D. Nigg, T. Monz, M. Chwalla, M. Hennrich, C. F. Roos, P. Zoller, and R. Blatt. “An open-system quantum simulator with trapped ions”. *Nature* **470**, 486–491 (2011). Cited on page 289.
- [Bar14] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O’Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and J. M. Martinis. “Superconducting quantum circuits at the surface code threshold for fault tolerance”. *Nature* **508**, 500–503 (2014). Cited on page 3.
- [Bau10] J. F. Bauters, M. J. Heck, D. John, M. C. Tien, A. Leinse, R. G. Heideman, D. J. Blumenthal, and J. E. Bowers. “Ultra-low loss silica-based waveguides with millimeter bend radius”. In: *Eur. Conf. Opt. Commun. ECOC*. Volume 1-2. 2010. ISBN: 9781424485352. Cited on page 135.
- [Bel64] J. S. Bell. “On the Einstein Podolsky Rosen paradox”. *Phys. Phys. Fiz.* **1**, 195–200 (1964). Cited on pages 4, 194.
- [Bel85] J. S. Bell. “The Theory of Local Beables”. *Dialectica* **39**, 86–96 (1985). Cited on page 197.
- [Ben02] C. H. Bennett. “Notes on Landauer’s principle, Reversible Computation and Maxwell’s Demon”. *Stud. Hist. Philos. Sci. Part B - Stud. Hist. Philos. Mod. Phys.* **34**, 501–510 (2002). Cited on page 287.
- [Ben06] I. Bengtsson and K. Życzkowski. *Geometry of Quantum States: An Introduction to Quantum Entanglement*. Cambridge University Press, 2006. Cited on page 187.
- [Ben08] J. Benhelm, G. Kirchmair, C. F. Roos, and R. Blatt. “Towards fault-tolerant quantum computing with trapped ions”. *Nat. Phys.* **4**, 463–466 (2008). Cited on page 2.

- [Ben14] C. D. Bentley, A. R. Carvalho, D. Kielpinski, and J. J. Hope. “Detection-enhanced steady state entanglement with ions”. *Phys. Rev. Lett.* **113**, 040501 (2014). Cited on page 288.
- [Ben73] C. H. Bennett. “Logical reversibility of computation”. *IBM J. Res. Dev.* **17**, 525–532 (1973). Cited on page 287.
- [Ben82] P. Benioff. “Quantum mechanical hamiltonian models of turing machines”. *J. Stat. Phys.* **29**, 515–546 (1982). Cited on page 1.
- [Ber21] M. Berlin-Udi, C. Matthiesen, P. N. T. Lloyd, A. Alonso, C. Noel, C. Orme, C.-E. Kim, V. Lordi, and H. Häffner. *Changes in electric-field noise due to thermal transformation of a surface ion trap*. 2021. arXiv: 2103.04482. Cited on pages 12, 332.
- [Ber97] E. Bernstein and U. Vazirani. “Quantum complexity theory”. *SIAM J. Comput.* **26**, 1411–1473 (1997). Cited on page 1.
- [Bir97] T. A. Birks, J. C. Knight, and P. S. J. Russell. “Endlessly single-mode photonic crystal fiber”. *Opt. Lett.* **22**, 961 (1997). Cited on pages 105, 106.
- [Bla08] R. Blatt and D. Wineland. *Entangled states of trapped atomic ions*. 2008. Cited on page 35.
- [Bla09] R. B. Blakestad, C. Ospelkaus, A. P. Vandevender, J. M. Amini, J. Britton, D. Leibfried, and D. J. Wineland. “High-fidelity transport of trapped-ion qubits through an X-junction trap array”. *Phys. Rev. Lett.* **102**, 153002 (2009). Cited on page 2.
- [Bla10] S. Blanes, F. Casas, J. A. Oteo, and J. Ros. “A pedagogical approach to the Magnus expansion”. *Eur. J. Phys.* **31**, 907–918 (2010). Cited on page 36.
- [Blo14] F. Blondelle, A. Sultan, E. Collin, H. Godfrin, F. Blondelle, A. Sultan, E. Collin, and H. Godfrin. “Electrical Conductance of Bolted Copper Joints for Cryogenic Applications”. *J. Low Temp. Phys.* 175 (2014). Cited on page 68.
- [Blu10] S. J. Blundell and K. M. Blundell. *Concepts in Thermal Physics*. Oxford University Press, 2010. ISBN: 9780191718236. Cited on page 72.
- [Boc18] M. Bock, P. Eich, S. Kucera, M. Kreis, A. Lenhard, C. Becher, and J. Eschner. “High-fidelity entanglement between a trapped ion and a telecom photon via quantum frequency conversion”. *Nat. Commun.* **9**, 1–7 (2018). Cited on page 106.
- [Bow12] R. Bowler, J. Gaebler, Y. Lin, T. R. Tan, D. Hanneke, J. D. Jost, J. P. Home, D. Leibfried, and D. J. Wineland. “Coherent Diabatic Ion Transport and Separation in a Multi-Zone Trap Array”. *Phys. Rev. Lett.* **109** (2012). Cited on page 2.

## BIBLIOGRAPHY

- [Bra16a] M. F. Brandl, M. W. van Mourik, L. Postler, A. Nolf, K. Lakhmanskiy, R. R. Paiva, S. Möller, N. Daniilidis, H. Häffner, V. Kaushal, T. Ruster, C. Warschburger, H. Kaufmann, U. G. Poschinger, F. Schmidt-Kaler, P. Schindler, T. Monz, and R. Blatt. “Cryogenic setup for trapped ion quantum computing”. *Rev. Sci. Instrum.* **87** (2016). Cited on pages 55, 68, 100, 332.
- [Bra16b] M. F. Brandl. “Towards Cryogenic Scalable Quantum Computing with Trapped Ions”. PhD thesis. University of Innsbruck, 2016. Cited on page 332.
- [Bra16c] M. F. Brandl, P. Schindler, T. Monz, and R. Blatt. “Cryogenic resonator design for trapped ion experiments in Paul traps”. *Appl. Phys. B Lasers Opt.* **122** (2016). Cited on page 83.
- [Bra18] S. Bravyi, D. Gosset, and R. König. “Quantum advantage with shallow circuits”. *Science (80-. )*. **362**, 308–311 (2018). Cited on page 4.
- [Bra19] C. E. Bradley, J. Randall, M. H. Abobeih, R. C. Berrevoets, M. J. Degen, M. A. Bakker, M. Markham, D. J. Twitchen, and T. H. Taminiau. “A Ten-Qubit Solid-State Spin Register with Quantum Memory up to One Minute”. *Phys. Rev. X* **9** (2019). Cited on page 3.
- [Bra20] S. Bravyi, D. Gosset, R. König, and M. Tomamichel. “Quantum advantage with noisy shallow circuits”. *Nat. Phys.* **16**, 1040–1045 (2020). Cited on page 4.
- [Bra90] S. L. Braunstein and C. M. Caves. “Wringing out better Bell inequalities”. *Ann. Phys. (N. Y.)*. **202**, 22–56 (1990). Cited on page 200.
- [Bro03] G. Brooker. *Modern Classical Optics*. Oxford University Press, 2003. Cited on page 97.
- [Bro11] K. R. Brown, A. C. Wilson, Y. Colombe, C. Ospelkaus, A. M. Meier, E. Knill, D. Leibfried, and D. J. Wineland. “Single-qubit-gate error below  $10^{-4}$  in a trapped ion”. *Phys. Rev. A - At. Mol. Opt. Phys.* **84**, 030303 (2011). Cited on page 67.
- [Bro14] M. Brownnutt, M. Kumph, P. Rabl, and R. Blatt. “Ion-trap measurements of electric-field noise near surfaces”. *Rev. Mod. Phys.* **87** (2014). Cited on pages 12, 55, 181.
- [Bru14] N. Brunner, D. Cavalcanti, S. Pironio, V. Scarani, and S. Wehner. “Bell nonlocality”. *Rev. Mod. Phys.* **86**, 419–478 (2014). Cited on page 220.
- [Bru15] C. D. Bruzewicz, J. M. Sage, and J. Chiaverini. “Measurement of ion motional heating rates over a range of trap frequencies and temperatures”. *Phys. Rev. A - At. Mol. Opt. Phys.* **91**, 041402 (2015). Cited on page 181.
- [Bru16] C. D. Bruzewicz, R. McConnell, J. Chiaverini, and J. M. Sage. “Scalable loading of a two-dimensional trapped-ion array”. *Nat. Commun.* **7**, 1–6 (2016). Cited on page 332.

- [Bru19] C. D. Bruzewicz, R. McConnell, J. Stuart, J. M. Sage, and J. Chiaverini. “Dual-species, multi-qubit logic primitives for  $\text{Ca}^+/\text{Sr}^+$  trapped-ion crystals”. *npj Quantum Inf.* **5**, 1–10 (2019). Cited on page 272.
- [Bud21] C. Budroni, A. Cabello, O. Gühne, M. Kleinmann, and J.-Å. Larsson. *Quantum Contextuality*. 2021. arXiv: 2102.13036. Cited on pages 4, 185, 224.
- [Bur10] A. H. Burrell. “High Fidelity Readout of Trapped Ion Qubits”. PhD thesis. University of Oxford, 2010. Cited on pages 33, 156.
- [Cab02] A. Cabello. “Finite-precision measurement does not nullify the Kochen-Specker theorem”. *Phys. Rev. A - At. Mol. Opt. Phys.* **65**, 4 (2002). Cited on page 221.
- [Cab09] A. Cabello. “Kochen-specker meets experiments”. *AIP Conf. Proc.* **1101**, 246–254 (2009). Cited on pages 203, 221, 222, 251.
- [Cab10a] A. Cabello and M. T. Cunha. “Proposal of a two-qutrit contextuality test free of the finite precision and compatibility loopholes”. *Phys. Rev. Lett.* **106** (2010). Cited on pages 221, 222.
- [Cab10b] A. Cabello, S. Severini, and A. Winter. *(Non-)Contextuality of Physical Theories as an Axiom*. 2010. arXiv: 1010.2163. Cited on pages 203, 204, 210, 212.
- [Cab13] A. Cabello. “Simple explanation of the quantum violation of a fundamental inequality”. *Phys. Rev. Lett.* **110**, 1–5 (2013). Cited on pages 210, 212.
- [Cal09] C. S. Calude and K. Svozil. *Quantum randomness and value indefiniteness*. 2009. Cited on page 228.
- [Cam14] E. T. Campbell. “Enhanced fault-tolerant quantum computing in  $d$ -level systems”. *Phys. Rev. Lett.* **113** (2014). Cited on page 334.
- [Cañ14a] G. Cañas, M. Arias, S. Etcheverry, E. S. Gómez, A. Cabello, G. B. Xavier, and G. Lima. “Applying the simplest kochen-specker set for quantum information processing”. *Phys. Rev. Lett.* **113**, 090404 (2014). Cited on page 239.
- [Cañ14b] G. Cañas, S. Etcheverry, E. S. Gómez, C. Saavedra, G. B. Xavier, G. Lima, and A. Cabello. “Experimental implementation of an eight-dimensional Kochen-Specker set and observation of its connection with the Greenberger-Horne-Zeilinger theorem”. *Phys. Rev. A - At. Mol. Opt. Phys.* **90**, 012119 (2014). Cited on page 239.
- [Cel08] S. Celozzi, R. Araneo, and G. Lovat. *Electromagnetic Shielding*. Wiley Series in Microwave and Optical Engineering. Hoboken, NJ, USA: John Wiley & Sons, Inc., 2008, pages 1–358. ISBN: 9780470268483. Cited on page 67.
- [Che20] J. -S. Chen, K. Wright, N. C. Pienti, D. Murphy, K. M. Beck, K. Landsman, J. M. Amini, and Y. Nam. “Efficient sideband cooling protocol for long trapped-ion chains”. *Phys. Rev. A* **102** (2020). Cited on page 52.

## BIBLIOGRAPHY

- [Chi05] J. Chiaverini, R. B. Blakestad, J. Britton, J. D. Jost, C. Langer, D. Leibfried, R. Ozeri, and D. J. Wineland. “Surface-Electrode Architecture for Ion-Trap Quantum Information Processing”. *Quantum Inf. Comput.* **5**, 419–439 (2005). Cited on pages 2, 11, 75, 77.
- [Chi14a] J. Chiaverini and J. M. Sage. “Insensitivity of the rate of ion motional heating to trap-electrode material over a large temperature range”. *Phys. Rev. A - At. Mol. Opt. Phys.* **89**, 012318 (2014). Cited on page 181.
- [Chi14b] G. Chiribella and X. Yuan. *Measurement sharpness cuts nonlocality and contextuality in every physical theory*. 2014. arXiv: 1404.3348. Cited on page 222.
- [Cho10] C. W. Chou, D. B. Hume, J. C. Koelemeij, D. J. Wineland, and T. Rosenband. “Frequency comparison of two high-accuracy Al<sup>+</sup> optical clocks”. *Phys. Rev. Lett.* **104**, 070802 (2010). Cited on page 152.
- [Cho11] J. Cho, S. Bose, and M. S. Kim. “Optical Pumping into Many-Body Entanglement”. *Phys. Rev. Lett.* **106**, 020504 (2011). Cited on page 288.
- [Chr10] M. Christiansen, R. Galea, D. Gong, S. Hou, D. Lissauer, C. Liu, T. Liu, V. Radeka, P. Rehak, J. Sondericker, R. Stroynowski, D. S. Su, P. Takacs, H. Takai, V. Tcherniatine, P. K. Teng, C. Thorn, A. C. Xiang, J. Ye, and B. Yu. “R&D towards cryogenic optical links”. *J. Instrum.* **5**, 20–24 (2010). Cited on page 143.
- [Chr15] B. G. Christensen, Y.-C. Liang, N. Brunner, N. Gisin, and P. G. Kwiat. “Exploring the Limits of Quantum Nonlocality with Entangled Photons”. *Phys. Rev. X* **5** (2015). Cited on pages 201, 256.
- [Cir80] B. S. Cirel’son. “Quantum generalizations of Bell’s inequality”. *Lett. Math. Phys.* **4**, 93–100 (1980). Cited on page 197.
- [Cir95] J. I. Cirac and P. Zoller. “Quantum computations with cold trapped ions”. *Phys. Rev. Lett.* **74**, 4091–4094 (1995). Cited on page 1.
- [Cla21] S. M. Clark, D. Lobser, M. Revelle, C. G. Yale, D. Bossert, A. D. Burch, M. N. Chow, C. W. Hogle, M. Ivory, J. Pehr, B. Salzbrenner, D. Stick, W. Sweatt, J. M. Wilson, E. Winrow, and P. Maunz. *Engineering the Quantum Scientific Computing Open User Testbed (QSCOUT): Design details and user guide*. 2021. arXiv: 2104.00759. Cited on pages 2, 259, 331.
- [Cla69] J. F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt. “Proposed experiment to test local hidden-variable theories”. *Phys. Rev. Lett.* **23**, 880–884 (1969). Cited on pages 194, 196.
- [Cla74] J. F. Clauser and M. A. Horne. “Experimental consequences of objective local theories”. *Phys. Rev. D* **10**, 526–535 (1974). Cited on page 197.

- [Cle04] R. Cleve, P. Høyer, B. Toner, and J. Watrous. “Consequences and limits of nonlocal strategies”. In: *Proc. Annu. IEEE Conf. Comput. Complex.* Volume 19. IEEE Computer Society, 2004, pages 236–249. Cited on pages 194, 196.
- [Cli00] R. Clifton. “Simulating quantum mechanics by non-contextual hidden variables”. *Proc. R. Soc. A Math. Phys. Eng. Sci.* **456**, 2101–2114 (2000). Cited on page 221.
- [Coh14] J. Cohen and M. Mirrahimi. “Dissipation-induced continuous quantum error correction for superconducting circuits”. *Phys. Rev. A* **90**, 062344 (2014). Cited on page 288.
- [Col09] R. Colbeck. “Quantum And Relativistic Protocols For Secure Multi-Party Computation”. PhD thesis. 2009. arXiv: 0911.3814. Cited on pages 198, 200.
- [Col11] R. Colbeck and R. Renner. “No extension of quantum theory can have improved predictive power”. *Nat. Commun.* **2** (2011). Cited on pages 201, 228.
- [Col21] D. C. Cole, J. J. Wu, S. D. Erickson, P.-Y. Hou, A. C. Wilson, D. Leibfried, and F. Reiter. “Dissipative preparation of W states in trapped ion systems”. (2021). Cited on pages 288, 338.
- [Con16] I. D. Conway Lamb. “Cryogenic Control Beyond 100 Qubits”. PhD thesis. University of Sydney, Masters Thesis, 2016. Cited on page 332.
- [Cor06] A. Corney. *Atomic and Laser Spectroscopy*. Oxford University Press, 2006. ISBN: 9780199211456. Cited on page 18.
- [Cro19] A. W. Cross, L. S. Bishop, S. Sheldon, P. D. Nation, and J. M. Gambetta. “Validating quantum computers using randomized model circuits”. *Phys. Rev. A* **100**, 032328 (2019). Cited on page 2.
- [DAm13] V. D’Ambrosio, I. Herbauts, E. Amsalem, E. Nagali, M. Bourennane, F. Sciarrino, and A. Cabello. “Experimental implementation of a Kochen-Specker set of quantum tests”. *Phys. Rev. X* **3**, 011012 (2013). Cited on page 239.
- [Den13] D. -.-L. Deng, C. Zu, X. -.-Y. Chang, P. -.-Y. Hou, H. -.-X. Yang, Y. -.-X. Wang, and L. -.-M. Duan. *Exploring Quantum Contextuality to Generate True Random Numbers*. 2013. arXiv: 1301.5364. Cited on page 252.
- [Des06] L. Deslauriers, S. Olmschenk, D. Stick, W. K. Hensinger, J. Sterk, and C. Monroe. “Scaling and suppression of anomalous heating in ion traps”. *Phys. Rev. Lett.* **97**, 103007 (2006). Cited on pages 12, 55.
- [Des19] B. Desiatov, A. Shams-Ansari, M. Zhang, C. Wang, and M. Lončar. “Ultra-low-loss integrated visible photonics using thin-film lithium niobate”. *Optica* **6**, 380 (2019). Cited on page 336.
- [Deu85] D. Deutsch. “Quantum theory, the Church–Turing principle and the universal quantum computer”. *Proc. R. Soc. London. A. Math. Phys. Sci.* **400**, 97–117 (1985). Cited on page 1.

## BIBLIOGRAPHY

- [Deu92] D. Deutsch and R. Jozsa. “Rapid solution of problems by quantum computation”. *Proc. R. Soc. London. Ser. A Math. Phys. Sci.* **439**, 553–558 (1992). Cited on page 1.
- [dNee20] B. de Neeve, T. L. Nguyen, T. Behrle, and J. Home. *Error correction of a logical grid state qubit by dissipative pumping*. 2020. arXiv: 2010.09681. Cited on pages 3, 288.
- [Dom08] C. Dominkovics and G. Harsányi. “Fractal description of dendrite growth during electrochemical migration”. *Microelectron. Reliab.* **48**, 1628–1634 (2008). Cited on page 79.
- [Don05] E. A. Donley, T. P. Heavner, F. Levi, M. O. Tataw, and S. R. Jefferts. “Double-pass acousto-optic modulator system”. *Rev. Sci. Instrum.* **76**, 063112 (2005). Cited on pages 102–104.
- [Don21] M. Dong, G. Clark, A. J. Leenheer, M. Zimmermann, D. Dominguez, A. J. Menssen, D. Heim, G. Gilbert, D. Englund, and M. Eichenfield. *High-speed programmable photonic circuits in a cryogenically compatible, visible-NIR 200 mm CMOS architecture*. 2021. arXiv: 2105.12531. Cited on page 337.
- [Dou20] E. Doucet, F. Reiter, L. Ranzani, and A. Kamal. “High fidelity dissipation engineering using parametric interactions”. *Phys. Rev. Res.* **2**, 023370 (2020). Cited on pages 288, 338.
- [Dre83] R. W. 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”. *Appl. Phys. B Photophysics Laser Chem.* **31**, 97–105 (1983). Cited on page 101.
- [Dut15] P. Duthil. “Material Properties at Low Temperature”. *CAS-CERN Accel. Sch. Supercond. Accel. - Proc.* 77–95 (2015). Cited on page 86.
- [Dzh12] E. N. Dzhafarov and J. V. Kujala. “All-Possible-Couplings Approach to Measuring Probabilistic Context”. *PLoS One* **8** (2012). Cited on page 222.
- [Eke91] A. K. Ekert. “Quantum cryptography based on Bell’s theorem”. *Phys. Rev. Lett.* **67**, 661–663 (1991). Cited on page 35.
- [Elt12] A. M. Eltony, S. X. Wang, G. M. Akselrod, P. F. Herskind, and I. L. Chuang. “Transparent ion trap with integrated photodetector”. *Appl. Phys. Lett.* **102** (2012). Cited on pages 134, 135, 336.
- [Esc03] J. Eschner, G. Morigi, F. Schmidt-Kaler, and R. Blatt. “Laser cooling of trapped ions”. *J. Opt. Soc. Am. B* **20**, 1003 (2003). Cited on page 49.
- [Fey82] R. P. Feynman. “Simulating physics with computers”. *Int. J. Theor. Phys.* **21**, 467–488 (1982). Cited on page 1.



- [Fey96] R. P. Feynman. *Feynman Lectures On Computation*. Addison-Wesley Publishing Company, Inc., 1996. ISBN: 0-201-48991-0. Cited on page 287.
- [Fin82] A. Fine. “Hidden Variables, Joint Probability, and the Bell Inequalities”. *Phys. Rev. Lett.* **48**, 291–295 (1982). Cited on pages 196, 197.
- [Fit15] J. F. Fitzsimons, J. A. Jones, and V. Vedral. “Quantum correlations which imply causation”. *Sci. Rep.* **5**, 18281 (2015). Cited on page 193.
- [Flü18] C. Flühmann, V. Negnevitsky, M. Marinelli, and J. P. Home. “Sequential Modular Position and Momentum Measurements of a Trapped Ion Mechanical Oscillator”. *Phys. Rev. X* **8**, 021001 (2018). Cited on pages 3, 22, 114.
- [Flü19] C. Flühmann. “Encoding a qubit in the motion of a trapped ion using superpositions of displaced squeezed states”. PhD thesis. ETH Zurich, 2019. Cited on pages 113, 125.
- [Fox03] R. W. Fox, C. W. Oates, and L. W. Hollberg. “1. Stabilizing diode lasers to high-finesse cavities”. *Exp. Methods Phys. Sci.* **40**, 1–46 (2003). Cited on pages 101, 113.
- [Fre72] S. J. Freedman and J. F. Clauser. “Experimental test of local hidden-variable theories”. *Phys. Rev. Lett.* **28**, 938–941 (1972). Cited on page 220.
- [Gae16] 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 Be 9 + Ion Qubits”. *Phys. Rev. Lett.* **117**, 060505 (2016). Cited on page 2.
- [Gan18] 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”. *Phys. Rev. A* **98** (2018). Cited on page 152.
- [Gao19] Y. Y. Gao, B. J. Lester, K. S. Chou, L. Frunzio, M. H. Devoret, L. Jiang, S. M. Girvin, and R. J. Schoelkopf. “Entanglement of bosonic modes through an engineered exchange interaction”. *Nature* **566**, 509–512 (2019). Cited on page 3.
- [Gar20] S. Garion, N. Kanazawa, H. Landa, D. C. McKay, S. Sheldon, A. W. Cross, and C. J. Wood. “Experimental implementation of non-Clifford interleaved randomized benchmarking with a controlled-S gate”. *arXiv* (2020). Cited on page 3.
- [Gea35] R. C. Geary. “The Ratio of the Mean Deviation to the Standard Deviation as a Test of Normality”. *Biometrika* **27**, 310 (1935). Cited on page 280.
- [Ger15] L. Gerster. *Spectral filtering and laser diode injection for multi-qubit trapped ion gates (MSc thesis)*. 2015. Cited on pages 114, 118.

## BIBLIOGRAPHY

- [Ger21] J. M. Gertler, B. Baker, J. Li, S. Shirol, J. Koch, and C. Wang. “Protecting a bosonic qubit with autonomous quantum error correction”. *Nature* **590**, 243–248 (2021). Cited on page 288.
- [Gha17] M. Ghadimi, V. Blums, 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”. *npj Quantum Inf.* **3**, 4 (2017). Cited on page 134.
- [Gok19] P. Gokhale, J. M. Baker, C. Duckering, N. C. Brown, K. R. Brown, and F. T. Chong. “Asymptotic Improvements to Quantum Circuits via Qutrits”. *Proc. - Int. Symp. Comput. Archit.* 554–566 (2019). Cited on page 334.
- [Got98] D. Gottesman. “Theory of fault-tolerant quantum computation”. *Phys. Rev. A - At. Mol. Opt. Phys.* **57**, 127–137 (1998). Cited on page 1.
- [Gre15] D. Greenbaum. *Introduction to Quantum Gate Set Tomography*. 2015. arXiv: 1509.02921. Cited on page 294.
- [Gri15] A. Grinbaum. “Quantum Theory as a Critical Regime of Language Dynamics”. *Found. Phys.* **45**, 1341–1350 (2015). Cited on page 228.
- [Gri17] D. J. Griffiths. *Introduction to Electrodynamics*. Cambridge University Press, 2017. Cited on page 91.
- [Gro96] L. K. Grover. “A fast quantum mechanical algorithm for database search”. In: *Proc. twenty-eighth Annu. ACM Symp. Theory Comput. - STOC '96*. Volume Part F1294. New York, New York, USA: ACM Press, 1996, pages 212–219. ISBN: 0897917855. arXiv: 9605043 [quant-ph]. Cited on page 1.
- [Gug15] M. Guggemos, D. Heinrich, O. A. Herrera-Sancho, R. Blatt, and C. F. Roos. “Sympathetic cooling and detection of a hot trapped ion by a cold one”. *New J. Phys.* **17**, 103001 (2015). Cited on page 148.
- [Gyg21] S. Gyger, J. Zichi, L. Schweickert, A. W. Elshaari, S. Steinhauer, S. F. Covre da Silva, A. Rastelli, V. Zwiller, K. D. Jöns, and C. Errando-Herranz. “Reconfigurable photonics with on-chip single-photon detectors”. *Nat. Commun.* **12**, 1–8 (2021). Cited on page 337.
- [Had86] G. R. Hadley. “Injection Locking of Diode Lasers”. *IEEE J. Quantum Electron.* **22**, 419–426 (1986). Cited on pages 117, 119.
- [Häf03] H. Häffner, S. Gulde, M. Riebe, G. Lancaster, C. Becher, J. Eschner, F. Schmidt-Kaler, and R. Blatt. “Precision Measurement and Compensation of Optical Stark Shifts for an Ion-Trap Quantum Processor”. *Phys. Rev. Lett.* **90**, 4 (2003). Cited on pages 274, 329.

- [Häf05] H. Häffner, W. Hänsel, C. F. Roos, J. Benhelm, D. Chek-al-kar, M. Chwalla, T. Körber, U. D. Rapol, M. Riebe, P. O. Schmidt, C. Becher, O. Gühne, W. Dür, and R. Blatt. “Scalable multiparticle entanglement of trapped ions”. *Nature* **438**, 643–646 (2005). Cited on page 338.
- [Häf08] H. Häffner, C. F. Roos, and R. Blatt. “Quantum computing with trapped ions”. *Phys. Rep.* **469**, 155–203 (2008). Cited on pages 7, 26, 35.
- [Hah50] E. L. Hahn. “Spin echoes”. *Phys. Rev.* **80**, 580–594 (1950). Cited on page 163.
- [Har10] M. Harlander, M. Brownnutt, W. Hänsel, and R. Blatt. “Trapped-ion probing of light-induced charging effects on dielectrics”. *New J. Phys.* **12** (2010). Cited on pages 12, 77, 269, 270.
- [Har14] 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”. *Phys. Rev. Lett.* **113** (2014). Cited on page 3.
- [Har19] K. Hartnett. *Quantum Supremacy Is Coming: Here’s What You Should Know*. 2019. (Visited on 06/16/2021). Cited on page 3.
- [Har96] S. Haroche and J. M. Raimond. “Quantum computing: dream or nightmare?”. *Phys. Today* **49**, 50–51 (1996). Cited on page 1.
- [Hay11] D. Hayes, K. Khodjasteh, L. Viola, and M. J. Biercuk. “Reducing sequencing complexity in dynamical quantum error suppression by Walsh modulation”. *Phys. Rev. A - At. Mol. Opt. Phys.* **84** (2011). Cited on pages 40, 284.
- [He19] Y. He, S. K. Gorman, D. Keith, L. Kranz, J. G. Keizer, and M. Y. Simmons. “A two-qubit gate between phosphorus donor electrons in silicon”. *Nature* **571**, 371–375 (2019). Cited on page 3.
- [Hec16] E. Hecht. *Optics*. Pearson, 2016. Cited on page 99.
- [Hem14] C. Hempel. “Digital quantum simulation, Schrodinger cat state spectroscopy and setting up a linear ion trap”. PhD thesis. University of Innsbruck, 2014. Cited on page 34.
- [Hen07] R. J. Hendricks, D. M. Grant, P. F. Herskind, A. Dantan, and M. Drewsen. “An all-optical ion-loading technique for scalable microtrap architectures”. *Appl. Phys. B Lasers Opt.* **88**, 507–513 (2007). Cited on page 149.
- [Hen15] B. Hensen, H. Bernien, A. E. Dreaú, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenbergh, R. F. 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). Cited on page 220.

## BIBLIOGRAPHY

- [Her17] M. Herrero-Collantes and J. C. Garcia-Escartin. “Quantum random number generators”. *Rev. Mod. Phys.* **89** (2017). Cited on page 228.
- [Her84] H. Hermann. *Waves and Fields in Optoelectronics*. q, 1984. Cited on page 136.
- [Hit13] D. A. Hite, Y. Colombe, A. C. Wilson, D. T. Allcock, D. Leibfried, D. J. Wineland, and D. P. Pappas. “Surface science for improved ion traps”. *MRS Bull.* **38**, 826–833 (2013). Cited on pages 2, 332.
- [Hit21] D. A. Hite, K. S. McKay, and D. P. Pappas. *Surface science motivated by heating of trapped ions from the quantum ground state*. 2021. arXiv: 2104.07818. Cited on page 332.
- [Hom06] J. Home. “Entanglement of Two Trapped-Ion Spin Qubits”. PhD thesis. University of Oxford, 2006. Cited on pages 43, 157, 174, 177, 260, 266.
- [Hom09a] J. P. Home, M. J. McDonnell, D. J. Szwer, B. C. Keitch, D. M. Lucas, D. N. Stacey, and A. M. Steane. “Memory coherence of a sympathetically cooled trapped-ion qubit”. *Phys. Rev. A - At. Mol. Opt. Phys.* **79**, 050305 (2009). Cited on page 2.
- [Hom09b] J. P. Home, D. Hanneke, J. D. Jost, J. M. Amini, D. Leibfried, and D. J. Wineland. “Complete methods set for scalable ion trap quantum information processing”. *Science (80-. )*. **325**, 1227–1230 (2009). Cited on page 2.
- [Hom13] J. P. Home. “Quantum science and metrology with mixed-species ion chains”. *Adv. At. Mol. Opt. Phys.* **62**, 231–277 (2013). Cited on page 264.
- [Hon16] S. Hong, M. Lee, H. Cheon, T. Kim, and D. I. D. Cho. “Guidelines for designing surface ion traps using the boundary element method”. *Sensors (Switzerland)* **16** (2016). Cited on page 12.
- [Hon21] Honeywell. *Honeywell Sets New Record For Quantum Computing Performance*. 2021. (Visited on 06/16/2021). Cited on page 2.
- [Hor10] K. Horodecki, M. Horodecki, P. Horodecki, R. Horodecki, M. Pawłowski, and M. Bourennane. *Contextuality offers device-independent security*. 2010. arXiv: 1006.0468v1. Cited on page 228.
- [Hor18] K. P. Horn, F. Reiter, Y. Lin, D. Leibfried, and C. P. Koch. “Quantum optimal control of the dissipative production of a maximally entangled state”. *New J. Phys.* **20**, 123010 (2018). Cited on page 288.
- [How14] M. Howard, J. Wallman, V. Veitch, and J. Emerson. “Contextuality supplies the ‘magic’ for quantum computation”. *Nature* **510**, 351–355 (2014). Cited on pages 4, 202.

- [Hu16] X. M. Hu, J. S. Chen, B. H. Liu, Y. Guo, Y. F. Huang, Z. Q. Zhou, Y. J. Han, C. F. Li, and G. C. Guo. “Experimental Test of Compatibility-Loophole-Free Contextuality with Spatially Separated Entangled Qutrits”. *Phys. Rev. Lett.* **117**, 170403 (2016). Cited on page 222.
- [Hua13] Y. F. Huang, M. Li, D. Y. Cao, C. Zhang, Y. S. Zhang, B. H. Liu, C. F. Li, and G. C. Guo. “Experimental test of state-independent quantum contextuality of an indivisible quantum system”. *Phys. Rev. A - At. Mol. Opt. Phys.* **87**, 052133 (2013). Cited on page 239.
- [Hua19] W. Huang, C. H. Yang, K. W. Chan, T. Tanttu, B. Hensen, R. C. Leon, M. A. Fogarty, J. C. Hwang, F. E. Hudson, K. M. Itoh, A. Morello, A. Laucht, and A. S. Dzurak. “Fidelity benchmarks for two-qubit gates in silicon”. *Nature* **569**, 532–536 (2019). Cited on page 3.
- [Hug10] I. Hughes and T. Hase. *Measurements and their Uncertainties : A practical guide to modern error analysis*. Oxford University Press, 2010. ISBN: 019956633X. Cited on page 277.
- [Hug11] M. D. Hughes, B. Lekitsch, J. A. Broersma, and W. K. Hensinger. “Microfabricated Ion Traps”. *Contemp. Phys.* **52**, 505–529 (2011). Cited on pages 12, 278.
- [Hug20] 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”. *Phys. Rev. Lett.* **125**, 080504 (2020). Cited on page 2.
- [Iba11] Y. Ibaraki, U. Tanaka, and S. Urabe. “Detection of parametric resonance of trapped ions for micromotion compensation”. *Appl. Phys. B Lasers Opt.* **105**, 219–223 (2011). Cited on page 147.
- [Iso] Isomet Corp. *AN0510 Acousto-Optic Modulation*. Technical report. Cited on pages 102, 104.
- [Ita] W. M. Itano, J. C. Bergquist, J. J. Bollinger, and D. J. Wineland. *Laser Cooling of Trapped Ions*. Technical report. Cited on page 52.
- [Iva06] P. A. Ivanov, E. S. Kyoseva, and N. V. Vitanov. “Engineering of arbitrary U(N) transformations by quantum Householder reflections”. *Phys. Rev. A - At. Mol. Opt. Phys.* **74**, 022323 (2006). Cited on page 234.
- [Ivo20] M. Ivory, W. J. Setzer, N. Karl, H. McGuinness, C. DeRose, M. Blain, D. Stick, M. Gehl, and L. P. Parazzoli. “Integrated optical addressing of a trapped ytterbium ion”. *arXiv* (2020). Cited on pages 134, 135, 270, 336.
- [Jam97] D. F. V. James. “Quantum dynamics of cold trapped ions, with application to quantum computation”. *Appl. Phys. B Lasers Opt.* **66**, 181–190 (1997). Cited on pages 16, 18, 99.

## BIBLIOGRAPHY

- [Jen00] T. Jennewein, U. Achleitner, G. Weihs, H. Weinfurter, and A. Zeilinger. “A fast and compact quantum random number generator”. *Rev. Sci. Instrum.* **71**, 1675–1680 (2000). Cited on page 229.
- [Jeo09] H. Jeong, M. Paternostro, and T. C. Ralph. “Failure of local realism revealed by extremely-coarse-grained measurements”. *Phys. Rev. Lett.* **102** (2009). Cited on page 256.
- [Jeo14] H. Jeong, Y. Lim, and M. S. Kim. “Coarsening measurement references and the quantum-to-classical transition”. *Phys. Rev. Lett.* **112** (2014). Cited on page 256.
- [Jer16] M. Jerger, Y. Reshitnyk, M. Oppliger, A. Potočník, M. Mondal, A. Wallraff, K. Goodenough, S. Wehner, K. Juliusson, N. K. Langford, and A. Fedorov. “Contextuality without nonlocality in a superconducting quantum system”. *Nat. Commun.* **7**, 7 (2016). Cited on pages 251, 252.
- [Joh16] K. G. Johnson, J. D. Wong-Campos, A. Restelli, K. A. Landsman, B. Neyenhuis, J. Mizrahi, and C. Monroe. “Active stabilization of ion trap radiofrequency potentials”. *Rev. Sci. Instrum.* **87**, 053110 (2016). Cited on page 83.
- [Jon12] J. A. Jones and D. Jaksch. *Quantum Information, Computation and Communication*. Cambridge: Cambridge University Press, 2012. ISBN: 9781139028509. Cited on page 188.
- [Jos20] M. K. Joshi, A. Fabre, C. Maier, T. Brydges, D. Kiesenhofer, H. Hainzer, R. Blatt, and C. F. Roos. “Polarization-gradient cooling of 1D and 2D ion Coulomb crystals”. *New J. Phys.* **22**, 103013 (2020). Cited on page 313.
- [Joz02] R. Jozsa and N. Linden. “On the role of entanglement in quantum computational speed-up”. *Proc. R. Soc. A Math. Phys. Eng. Sci.* **459**, 2011–2032 (2002). Cited on page 4.
- [Kas11] M. J. Kastoryano, F. Reiter, and A. S. Sørensen. “Dissipative Preparation of Entanglement in Optical Cavities”. *Phys. Rev. Lett.* **106**, 090502 (2011). Cited on page 288.
- [Kau20] V. Kaushal, B. Lekitsch, A. Stahl, J. Hilder, D. Pijn, C. Schmiegelow, A. Bermudez, M. Müller, F. Schmidt-Kaler, and U. Poschinger. “Shuttling-based trapped-ion quantum information processing”. *AVS Quantum Sci.* **2**, 014101 (2020). Cited on page 258.
- [Kei07] B. Keitch. “A Quantum Memory Qubit in Calcium-43”. PhD thesis. University of Oxford, 2007. Cited on page 116.
- [Kel15] J. Keller, H. L. Partner, T. Burgermeister, and T. E. Mehlstäubler. “Precise determination of micromotion for trapped-ion optical clocks”. *J. Appl. Phys.* **118**, 104501 (2015). Cited on pages 10, 146.

- [Ken99] A. Kent. “Noncontextual hidden variables and physical measurements”. *Phys. Rev. Lett.* **83**, 3755–3757 (1999). Cited on page 221.
- [Kes11] A. Keselman, Y. Glickman, N. Akerman, S. Kotler, and R. Ozeri. “High-fidelity state detection and tomography of a single-ion Zeeman qubit”. *New J. Phys.* **13**, 073027 (2011). Cited on page 18.
- [Kha15] S. Khan, M. P. Kumar, V. Bharti, and V. Natarajan. “Coherent population trapping (CPT) versus electromagnetically induced transparency (EIT)”. *Eur. Phys. J. D* **71** (2015). Cited on page 50.
- [Kie00a] D. Kielpinski, B. E. King, C. J. Myatt, C. A. Sackett, Q. A. Turchette, W. M. Itano, C. Monroe, D. J. Wineland, and W. H. Zurek. “Sympathetic cooling of trapped ions for quantum logic”. *Phys. Rev. A - At. Mol. Opt. Phys.* **61**, 8 (2000). Cited on page 2.
- [Kie00b] D. Kielpinski, B. E. King, C. J. Myatt, C. A. Sackett, Q. A. Turchette, W. M. Itano, C. Monroe, D. J. Wineland, and W. H. Zurek. “Sympathetic cooling of trapped ions for quantum logic”. *Phys. Rev. A - At. Mol. Opt. Phys.* **61**, 8 (2000). Cited on page 330.
- [Kie02] D. Kielpinski, C. Monroe, and D. J. Wineland. “Architecture for a large-scale ion-trap quantum computer”. *Nature* **417**, 709–711 (2002). Cited on page 2.
- [Kie15] D. Kienzler, B. Keitch, L. D. Clercq, F. Leupold, F. Lindenefser, M. Marinelli, V. Negnevitsky, and J. P. Home. “Quantum harmonic oscillator state synthesis by reservoir engineering”. *Science (80-. )*. **347**, 53–56 (2015). Cited on page 288.
- [Kie17] D. Kienzler, H. Y. Lo, V. Negnevitsky, C. Flühmann, M. Marinelli, and J. P. Home. “Quantum Harmonic Oscillator State Control in a Squeezed Fock Basis”. *Phys. Rev. Lett.* **119**, 033602 (2017). Cited on page 3.
- [Kim07] J. Kim and C. Kim. “Integrated Optical Approach to Trapped Ion Quantum Computation”. *Quantum Inf. Comput.* **9**, 181–202 (2007). Cited on page 134.
- [Kim10] K. Kim, M. S. Chang, S. Korenblit, R. Islam, E. E. Edwards, J. K. Freericks, G. D. Lin, L. M. Duan, and C. Monroe. “Quantum simulation of frustrated Ising spins with trapped ions”. *Nature* **465**, 590–593 (2010). Cited on page 295.
- [Kim11] T. H. Kim, P. F. Herskind, and I. L. Chuang. “Surface-electrode ion trap with integrated light source”. *Appl. Phys. Lett.* **98**, 214103 (2011). Cited on pages 134, 135.
- [Kim16] M. E. Kimchi-Schwartz, L. Martin, E. Flurin, C. Aron, M. Kulkarni, H. E. Tureci, and I. Siddiqi. “Stabilizing Entanglement via Symmetry-Selective Bath Engineering in Superconducting Qubits”. *Phys. Rev. Lett.* **116**, 240503 (2016). Cited on page 288.

## BIBLIOGRAPHY

- [Kin98] B. E. King, C. S. Wood, C. J. Myatt, Q. A. Turchette, D. Leibfried, W. M. Itano, C. Monroe, and D. J. Wineland. “Cooling the collective motion of trapped ions to initialize a quantum register”. *Phys. Rev. Lett.* **81**, 1525–1528 (1998). Cited on pages 1, 177, 260, 264.
- [Kir09a] G. Kirchmair, J. Benhelm, F. Zähringer, R. Gerritsma, C. F. Roos, and R. Blatt. “Deterministic entanglement of ions in thermal states of motion”. *New J. Phys.* **11**, 023002 (2009). Cited on page 274.
- [Kir09b] G. Kirchmair, F. Zähringer, R. Gerritsma, M. Kleinmann, O. Gühne, A. Cabello, R. Blatt, and C. F. Roos. “State-independent experimental test of quantum contextuality”. *Nature* **460**, 494–497 (2009). Cited on pages 225, 239.
- [Kja20] M. Kjaergaard, M. E. Schwartz, A. Greene, G. O. Samach, A. Bengtsson, M. O’Keeffe, C. M. McNally, J. Braumüller, D. K. Kim, P. Krantz, M. Marvian, A. Melville, B. M. Niedzielski, Y. Sung, R. Winik, J. Yoder, D. Rosenberg, K. Obenland, S. Lloyd, T. P. Orlando, I. Marvian, S. Gustavsson, and W. D. Oliver. “Programming a quantum computer with quantum instructions”. *arXiv* (2020). Cited on page 3.
- [Kle12] M. Kleinmann, C. Budroni, J. Å. Larsson, O. Gühne, and A. Cabello. “Optimal inequalities for state-independent contextuality”. *Phys. Rev. Lett.* **109**, 1–5 (2012). Cited on page 219.
- [Kly08] A. A. Klyachko, M. A. Can, S. Binicioğlu, and A. S. Shumovsky. “Simple test for hidden variables in spin-1 systems”. *Phys. Rev. Lett.* **101**, 1–4 (2008). Cited on pages 204, 206.
- [Kni96] J. C. Knight, T. A. Birks, P. S. J. Russell, and D. M. Atkin. “All-silica single-mode optical fiber with photonic crystal cladding”. *Opt. Lett.* **21**, 1547 (1996). Cited on page 106.
- [Kof07] J. Kofler and Č. Brukner. “Classical world arising out of quantum physics under the restriction of coarse-grained measurements”. *Phys. Rev. Lett.* **99** (2007). Cited on page 256.
- [Kon16] Kondephy. *Pound–Drever–Hall technique.svg* - *Wikimedia Commons*. 2016. (Visited on 06/02/2021). Cited on page 101.
- [Kra08] B. Kraus, H. P. Büchler, S. Diehl, A. Kantian, A. Micheli, and P. Zoller. “Preparation of entangled states by quantum Markov processes”. *Phys. Rev. A* **78**, 042307 (2008). Cited on page 288.
- [Kri18] S. Krinner, S. Storz, P. Kurpiers, P. Magnard, J. Heinsoo, R. Keller, J. Luetolf, C. Eichler, and A. Wallraff. “Engineering cryogenic setups for 100-qubit scale superconducting circuit systems”. *EPJ Quantum Technol.* **6** (2018). Cited on page 86.



- [Kuj15] J. V. Kujala, E. N. Dzhafarov, and J. Å. Larsson. “Necessary and Sufficient Conditions for an Extended Noncontextuality in a Broad Class of Quantum Mechanical Systems”. *Phys. Rev. Lett.* **115**, 150401 (2015). Cited on pages 222, 227, 246.
- [Kul17] A. Kulikov, M. Jerger, A. Potočník, A. Wallraff, and A. Fedorov. “Realization of a Quantum Random Generator Certified with the Kochen-Specker Theorem”. *Phys. Rev. Lett.* **119** (2017). Cited on page 228.
- [Kun15] R. Kunjwal and R. W. Spekkens. “From the Kochen-Specker Theorem to Noncontextuality Inequalities without Assuming Determinism”. *Phys. Rev. Lett.* **115** (2015). Cited on page 223.
- [Kun19] R. Kunjwal. “Beyond the Cabello-Severini-Winter framework: Making sense of contextuality without sharpness of measurements”. *Quantum* **3** (2019). Cited on pages 222, 223, 253.
- [Lab07] J. Labaziewicz, Y. Ge, P. Antohi, D. Leibbrandt, K. R. Brown, and I. L. Chuang. “Suppression of Heating Rates in Cryogenic Surface-Electrode Ion Traps”. *Phys. Rev. Lett.* **100** (2007). Cited on pages 2, 55.
- [Lan09] B. P. Lanyon, M. Barbieri, M. P. Almeida, T. Jennewein, T. C. Ralph, K. J. Resch, G. J. Pryde, J. L. O’Brien, A. Gilchrist, and A. G. White. “Simplifying quantum logic using higher-dimensional Hilbert spaces”. *Nat. Phys.* **5**, 134–140 (2009). Cited on page 334.
- [Lan61] R. Landauer. “Irreversibility and Heat Generation in the Computing Process”. *IBM J. Res. Dev.* **5**, 183–191 (1961). Cited on page 287.
- [Lap11] R. Lapkiewicz, P. Li, C. Schaeff, N. K. Langford, S. Ramelow, M. Wieśniak, and A. Zeilinger. “Experimental non-classicality of an indivisible quantum system”. *Nature* **474**, 490–493 (2011). Cited on page 252.
- [Lap13] R. Lapkiewicz, P. Li, C. Schaeff, N. K. Langford, S. Ramelow, M. Wiesniak, and A. Zeilinger. *Comment on "Two Fundamental Experimental Tests of Non-classicality with Qutrits"*. 2013. arXiv: 1305.5529. Cited on page 253.
- [Lar02] J. A. Larsson. “A Kochen-Specker inequality”. *Europhys. Lett.* **58**, 799–805 (2002). Cited on page 222.
- [Lar19] M. V. Larsen, X. Guo, C. R. Breum, J. S. Neergaard-Nielsen, and U. L. Andersen. “Deterministic generation of a two-dimensional cluster state”. *Science (80-. )*. **366**, 369–372 (2019). Cited on page 3.
- [Lec10] R. Lechner. “Photoionisation of 40 Ca with a frequency-doubled 422 nm laser and a 377 nm laser diode (MSc thesis)”. PhD thesis. University of Innsbruck, 2010. Cited on page 149.

## BIBLIOGRAPHY

- [Lec16] R. Lechner, C. Maier, C. Hempel, P. Jurcevic, B. P. Lanyon, T. Monz, M. Brownnutt, R. Blatt, and C. F. Roos. “Electromagnetically-induced-transparency ground-state cooling of long ion strings”. *Phys. Rev. A* **93**, 053401 (2016). Cited on page 51.
- [Lee00] K. K. Lee, D. R. Lim, H. C. Luan, A. Agarwal, J. Foresi, and L. C. Kimerling. “Effect of size and roughness on light transmission in a Si/SiO<sub>2</sub> waveguide: Experiments and model”. *Appl. Phys. Lett.* **77**, 1617–1619 (2000). Cited on page 135.
- [Lei03a] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland. “Quantum dynamics of single trapped ions”. *Rev. Mod. Phys.* **75**, 281–324 (2003). Cited on pages 7, 14, 19.
- [Lei03b] D. Leibfried, B. DeMarco, V. Meyer, D. Lucas, M. Barrett, J. Britton, W. M. Itano, B. Jelenković, C. Langer, T. Rosenband, and D. J. Wineland. “Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate”. *Nature* **422**, 412–415 (2003). Cited on page 2.
- [Lei07] D. R. Leibbrandt, R. J. Clark, J. Labaziewicz, P. Antohi, W. Bakr, K. R. Brown, and I. L. Chuang. “Laser ablation loading of a surface-electrode ion trap”. *Phys. Rev. A - At. Mol. Opt. Phys.* **76** (2007). Cited on pages 13, 149.
- [Lei09] D. R. Leibbrandt, J. Labaziewicz, R. J. Clark, I. L. Chuang, R. J. Epstein, C. Ospelkaus, J. H. Wesenberg, J. J. Bollinger, D. Leibfried, D. J. Wineland, D. Stick, J. Sterk, C. Monroe, C. -.-S. Pai, Y. Low, R. Frahm, and R. E. Slusher. “Demonstration of a scalable, multiplexed ion trap for quantum information processing”. *Quantum Inf. Comput.* **9**, 0901–0919 (2009). Cited on page 12.
- [Lei11] D. R. Leibbrandt, M. J. Thorpe, M. Notcutt, R. E. Drullinger, T. Rosenband, and J. C. Bergquist. “Spherical reference cavities for frequency stabilization of lasers in non-laboratory environments”. *Opt. Express* **19**, 3471 (2011). Cited on page 113.
- [Leo19] T. Leopold, S. A. King, P. Micke, A. Bautista-Salvador, J. C. Heip, C. Ospelkaus, J. R. Crespo López-Urrutia, and P. O. Schmidt. “A cryogenic radio-frequency ion trap for quantum logic spectroscopy of highly charged ions”. *Rev. Sci. Instrum.* **90**, 073201 (2019). Cited on page 148.
- [Leo61] F. C. Leone, L. S. Nelson, and R. B. Nottingham. “The Folded Normal Distribution”. *Technometrics* **3**, 543 (1961). Cited on page 249.
- [Leu15] F. Leupold. “Bang-bang Control of a Trapped-Ion Oscillator”. PhD thesis. 2015. Cited on pages 18, 53, 68, 70, 75, 76, 82, 89, 104, 131, 156.
- [Leu17] F. M. Leupold, M. Malinowski, C. Zhang, V. Negnevitsky, J. Alonso, A. Cabello, and J. P. Home. “Sustained state-independent quantum contextual correlations from a single ion”. *Phys. Rev. Lett.* **120** (2017). Cited on pages 217, 244.

- [Lev18] H. Levine, A. Keesling, A. Omran, H. Bernien, S. Schwartz, A. S. Zibrov, M. Endres, M. Greiner, V. Vuletić, and M. D. Lukin. “High-Fidelity Control and Entanglement of Rydberg-Atom Qubits”. *Phys. Rev. Lett.* **121**, 123603 (2018). Cited on page 3.
- [Lin11] F. Lindenefelser. *Laser stabilisation for quantum information experiments with trapped ions (MSc thesis)*. 2011. Cited on page 101.
- [Lin13] Y. Lin, J. P. Gaebler, F. Reiter, T. R. Tan, R. Bowler, A. S. Sørensen, D. Leibfried, and D. J. Wineland. “Dissipative production of a maximally entangled steady state of two quantum bits”. *Nature* **504**, 415–418 (2013). Cited on page 288.
- [Lin16] Y. Lin, J. P. Gaebler, F. Reiter, T. R. Tan, R. Bowler, Y. Wan, A. Keith, E. Knill, S. Glancy, K. Coakley, A. S. Sørensen, D. Leibfried, and D. J. Wineland. “Preparation of Entangled States through Hilbert Space Engineering”. *Phys. Rev. Lett.* **117**, 140502 (2016). Cited on page 288.
- [Lin20] A. Lindberg. *Improving the coherent quantum control of trapped ion qubits (MSc thesis)*. Technical report. Stockholm University, 2020. Cited on page 169.
- [Lin76] G. Lindblad. “On the generators of quantum dynamical semigroups”. *Commun. Math. Phys.* **48**, 119–130 (1976). Cited on page 24.
- [Liu21] H. Liu, G. H. Low, D. S. Steiger, T. Häner, M. Reiher, and M. Troyer. *Prospects of Quantum Computing for Molecular Sciences*. 2021. arXiv: 2102.10081. Cited on page 3.
- [Lo15] H. Y. Lo, D. Kienzler, L. De Clercq, M. Marinelli, V. Negnevitsky, B. C. Keitch, and J. P. Home. “Spin-motion entanglement and state diagnosis with squeezed oscillator wavepackets”. *Nature* **521**, 336–339 (2015). Cited on page 288.
- [Low19] P. J. Low, B. M. White, A. A. Cox, M. L. Day, and C. Senko. “Practical trapped-ion protocols for universal qudit-based quantum computing”. *arXiv* (2019). Cited on pages 232, 334.
- [Luc03] D. M. Lucas, A. Ramos, J. P. Home, M. J. McDonnell, S. Nakayama, J. .-.P. Stacey, S. C. Webster, D. N. Stacey, and A. M. Steane. “Isotope-selective photo-ionization for calcium ion trapping”. *Phys. Rev. A - At. Mol. Opt. Phys.* **69**, 13 (2003). Cited on pages 13, 72.
- [Lud98] H. Ludvigsen, M. Tossavainen, and M. Kaivola. “Laser linewidth measurements using self-homodyne detection with short delay”. *Opt. Commun.* **155**, 180–186 (1998). Cited on page 124.
- [Ma94] L.-S. Ma, P. Jungner, J. Ye, and J. L. Hall. “Delivering the same optical frequency at two places: accurate cancellation of phase noise introduced by an optical fiber or other time-varying path”. *Opt. Lett.* **19**, 1777 (1994). Cited on page 114.

## BIBLIOGRAPHY

- [Mad20] I. S. Madjarov, J. P. Covey, A. L. Shaw, J. Choi, A. Kale, A. Cooper, H. Pichler, V. Schkolnik, J. R. Williams, and M. Endres. “High-fidelity entanglement and detection of alkaline-earth Rydberg atoms”. *Nat. Phys.* **16**, 857–861 (2020). Cited on page 3.
- [Mal17] M. Malinowski, C. Zhang, F. M. Leupold, A. Cabello, J. Alonso, and J. P. Home. “Probing the limits of correlations in an indivisible quantum system”. *Phys. Rev. A* **98** (2017). Cited on pages 190, 207, 255.
- [Mal21] M. Malinowski, C. Zhang, V. Negnevitsky, I. Rojkov, F. Reiter, M. Stadler, D. Kienzler, K. K. Mehta, and J. P. Home. *(In preparation) Generation of a maximally entangled state using collective optical pumping*. 2021. Cited on pages 289, 296.
- [Mar17] C. D. Marciniak, H. B. Ball, A. T. .-.H. Hung, and M. J. Biercuk. “Towards fully commercial, UV-compatible fiber patch cords”. *arXiv* **453**, 1–86 (2017). Cited on page 129.
- [Mar20] M. Marinelli. “Quantum information processing with mixed-species ion crystals”. PhD thesis. ETH Zurich, 2020. Cited on pages 34, 93, 158, 271.
- [Mas06] L. Masanes, A. Acin, and N. Gisin. “General properties of nonsignaling theories”. *Phys. Rev. A - At. Mol. Opt. Phys.* **73** (2006). Cited on page 198.
- [Maz11] J. R. Maze, A. Gali, E. Togan, Y. Chu, A. Trifonov, E. Kaxiras, and M. D. Lukin. “Properties of nitrogen-vacancy centers in diamond: the group theoretic approach”. *New J. Phys.* **13**, 025025 (2011). Cited on page 338.
- [Maz15] M. D. Mazurek, M. F. Pusey, R. Kunjwal, K. J. Resch, and R. W. Spekkens. “An experimental test of noncontextuality without unwarranted idealizations”. *Nat. Commun.* **7**, 1–19 (2015). Cited on page 223.
- [Maz17] M. D. Mazurek, M. F. Pusey, K. J. Resch, and R. W. Spekkens. “Experimentally bounding deviations from quantum theory in the landscape of generalized probabilistic theories”. *PRX Quantum* **2** (2017). Cited on page 335.
- [Maz18] E. Mazzeo. *The challenge of filtering photons in a cold environment*. Technical report. DESY, 2018. Cited on page 143.
- [McK16] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta. “Efficient Z-Gates for Quantum Computing”. *Phys. Rev. A* **96** (2016). Cited on page 26.
- [McL64] N. W. McLachlan. *Theory and Application of Mathieu Functions*. Dover Publications, 1964. ISBN: 9780486612331. Cited on page 9.
- [Mee96] D. M. Meekhof, C. Monroe, B. E. King, W. M. Itano, and D. J. Wineland. “Generation of nonclassical motional states of a trapped atom”. *Phys. Rev. Lett.* **76**, 1796–1799 (1996). Cited on page 177.

- [Meh14] K. K. Mehta, A. M. Eltony, C. D. Bruzewicz, I. L. Chuang, R. J. Ram, J. M. Sage, and J. Chiaverini. “Ion traps fabricated in a CMOS foundry”. *Appl. Phys. Lett.* **105** (2014). Cited on pages 13, 76.
- [Meh16] K. K. Mehta, C. D. Bruzewicz, R. McConnell, R. J. Ram, J. M. Sage, and J. Chiaverini. “Integrated optical addressing of an ion qubit”. *Nat. Nanotechnol.* **11**, 1066–1070 (2016). Cited on pages 13, 134, 135.
- [Meh17] K. K. Mehta and R. J. Ram. “Precise and diffraction-limited waveguide-to-free-space focusing gratings”. *Sci. Rep.* **7**, 1–8 (2017). Cited on pages 140, 142.
- [Meh20] 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). Cited on pages 2, 82, 100, 134, 137, 139, 257.
- [Mer11] J. T. Merrill, C. Volin, D. Landgren, J. M. Amini, K. Wright, S. C. Doret, C. S. Pai, H. Hayden, T. Killian, D. Faircloth, K. R. Brown, A. W. Harter, and R. E. Slusher. “Demonstration of integrated microscale optics in surface-electrode ion traps”. *New J. Phys.* **13**, 103005 (2011). Cited on page 134.
- [Mer93] N. D. Mermin. “Hidden Variables and the Two Theorems of John Bell”. *Rev. Mod. Phys.* **65**, 803–815 (1993). Cited on pages 204, 223–225.
- [Mey99] D. A. Meyer. “Finite precision measurement nullifies the kochen-specker theorem”. *Phys. Rev. Lett.* **83**, 3751–3754 (1999). Cited on page 221.
- [Mic19] P. Micke, J. Stark, S. A. King, T. Leopold, T. Pfeifer, L. Schmöger, M. Schwarz, L. J. Spieß, P. O. Schmidt, and J. R. Crespo López-Urrutia. “Closed-cycle, low-vibration 4 K cryostat for ion traps and other applications”. *Rev. Sci. Instrum.* **90**, 65104 (2019). Cited on page 55.
- [Mil18a] A. Militaru. *An Optical Setup for Single Ion Addressing in a Cryogenic Surface-Electrode Ion Trap (MSc thesis)*. 2018. Cited on page 100.
- [Mil18b] A. R. Milne, C. L. Edmunds, C. Hempel, F. Roy, S. Mavadia, and M. J. Biercuk. “Phase-modulated entangling gates robust to static and time-varying errors”. *Phys. Rev. Appl.* **13** (2018). Cited on pages 42, 305.
- [Mil19] S. Miller. “Optical pulse shaping for trapped-ion quantum computing with integrated photonics”. PhD thesis. Masters Thesis, ETH Zurich, 2019. Cited on page 137.
- [Mit20] A. Mitra, M. J. Martin, G. W. Biedermann, A. M. Marino, P. M. Poggi, and I. H. Deutsch. “Robust Mølmer-Sørensen gate for neutral atoms using rapid adiabatic Rydberg dressing”. *Phys. Rev. A* **101**, 030301 (2020). Cited on page 338.

## BIBLIOGRAPHY

- [Mon09] T. Monz, K. Kim, A. S. Villar, P. Schindler, M. Chwalla, M. Riebe, C. F. Roos, H. Häffner, W. Hänsel, M. Hennrich, and R. Blatt. “Realization of Universal Ion Trap Quantum Computation with Decoherence Free Qubits”. *Phys. Rev. Lett.* **103** (2009). Cited on page 334.
- [Mon21] C. Monroe, W. C. Campbell, L. M. Duan, Z. X. Gong, A. V. Gorshkov, P. W. Hess, R. Islam, K. Kim, N. M. Linke, G. Pagano, P. Richerme, C. Senko, and N. Y. Yao. “Programmable quantum simulations of spin systems with trapped ions”. *Rev. Mod. Phys.* **93**, 025001 (2021). Cited on page 2.
- [Mon95] C. Monroe, D. M. Meekhof, B. E. King, W. M. Itano, and D. J. Wineland. “Demonstration of a fundamental quantum logic gate”. *Phys. Rev. Lett.* **75**, 4714–4717 (1995). Cited on page 1.
- [Mor00] G. Morigi, J. Eschner, and C. H. Keitel. “Ground state laser cooling using electromagnetically induced transparency”. *Phys. Rev. Lett.* **85**, 4458–4461 (2000). Cited on pages 49, 50.
- [Mor15] G. Morigi, J. Eschner, C. Cormick, Y. Lin, D. Leibfried, and D. J. Wineland. “Dissipative Quantum Control of a Spin Chain”. *Phys. Rev. Lett.* **115**, 200502 (2015). Cited on page 288.
- [Mor99] G. Morigi, J. Eschner, J. I. Cirac, and P. Zoller. “Laser cooling of two trapped ions: Sideband cooling beyond the Lamb-Dicke limit”. *Phys. Rev. A - At. Mol. Opt. Phys.* **59**, 3797–3808 (1999). Cited on page 52.
- [Mou10] O. Moussa, C. A. Ryan, D. G. Cory, and R. Laflamme. “Testing Contextuality on Quantum Ensembles with One Clean Qubit”. *Phys. Rev. Lett.* **104**, 160501 (2010). Cited on page 239.
- [Mye08] A. H. Myerson, D. J. Szwer, S. C. Webster, D. T. 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”. *Phys. Rev. Lett.* **100**, 200502 (2008). Cited on page 32.
- [Näg99] H. C. Nägerl, D. Leibfried, H. Rohde, G. Thalhammer, J. Eschner, F. Schmidt-Kaler, and R. Blatt. “Laser addressing of individual ions in a linear ion trap”. *Phys. Rev. A - At. Mol. Opt. Phys.* **60**, 145–148 (1999). Cited on page 1.
- [Neg18a] V. Negnevitsky, M. Marinelli, K. K. Mehta, H. Y. Lo, C. Flühmann, and J. P. Home. “Repeated multi-qubit readout and feedback with a mixed-species trapped-ion register”. *Nature* **563**, 527–531 (2018). Cited on pages 158, 236.
- [Neg18b] V. Negnevitsky. “Feedback-stabilised quantum states in a mixed-species ion system”. PhD thesis. ETH Zurich, 2018. Cited on pages 92, 93, 296.
- [Nie02] M. A. Nielsen and I. Chuang. *Quantum computation and quantum information*. American Association of Physics Teachers, 2002. Cited on pages 187, 189.

- [Nie08] X. R. Nie, C. F. Roos, and D. F. V. James. “Theory of Cross Phase Modulation for the Vibrational Modes of Trapped Ions”. *Phys. Lett. Sect. A Gen. At. Solid State Phys.* **373**, 422–425 (2008). Cited on page 47.
- [Nif20] R. J. Niffenegger, J. Stuart, C. Sorace-Agaskar, D. Kharas, S. Bramhavar, C. D. Bruzewicz, W. Loh, R. T. Maxson, R. McConnell, D. Reens, G. N. West, J. M. Sage, and J. Chiaverini. “Integrated multi-wavelength control of an ion qubit”. *Nature* **586**, 538–542 (2020). Cited on pages 2, 134, 135, 270, 335, 336.
- [Nod86] J. Noda, K. Okamoto, and Y. Sasaki. “Polarization-Maintaining Fibers and Their Applications”. *J. Light. Technol.* **4**, 1071–1089 (1986). Cited on page 106.
- [Noe19] C. Noel. “High temperature studies of electric-field noise in a surface ion trap”. PhD thesis. University of California Berkeley, 2019. Cited on page 178.
- [OLe06] D. P. O’Leary, G. K. Brennen, and S. S. Bullock. “Parallelism for Quantum Computation with Qudits”. *Phys. Rev. A - At. Mol. Opt. Phys.* **74** (2006). Cited on page 233.
- [OMa15] P. J. O’Malley, J. Kelly, R. Barends, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, I. C. Hoi, E. Jeffrey, A. Megrant, J. Mutus, C. Neill, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, A. N. Korotkov, A. N. Cleland, and J. M. Martinis. “Qubit Metrology of Ultralow Phase Noise Using Randomized Benchmarking”. *Phys. Rev. Appl.* **3** (2015). Cited on page 164.
- [Opp15] N. Oppong. *Cancellation of optical phase noise induced by an optical fiber*. Technical report. Project report, ETH Zurich, 2015. Cited on page 114.
- [Osp11] C. Ospelkaus, U. Warring, Y. Colombe, K. R. Brown, J. M. Amini, D. Leibfried, and D. J. Wineland. *Microwave quantum logic gates for trapped ions*. 2011. arXiv: 1104.3573. Cited on pages 13, 337.
- [Oze11] R. Ozeri. “Tutorial: The trapped-ion qubit toolbox”. *Contemp. Phys.* **52**, 531–550 (2011). Cited on pages 18, 22.
- [Pag19] G. Pagano, P. W. Hess, H. B. Kaplan, W. L. Tan, P. Richerme, P. Becker, A. Kyprianidis, J. Zhang, E. Birkelbaw, M. R. Hernandez, Y. Wu, and C. Monroe. “Cryogenic trapped-ion system for large scale quantum simulation”. *Quantum Sci. Technol.* **4**, 14004 (2019). Cited on pages 55, 71.
- [Pap16] M. Papes, P. Cheben, D. Benedikovic, J. H. Schmid, J. Pond, R. Halir, A. Ortega-Moñux, G. Wangüemert-Pérez, W. N. Ye, D.-X. Xu, S. Janz, M. Dado, and V. Vašínek. “Fiber-chip edge coupler with large mode size for silicon photonic wire waveguides”. *Opt. Express* **24**, 5026 (2016). Cited on page 138.
- [Pea05] C. E. Pearson, D. R. Leibbrandt, W. S. Bakr, W. J. Mallard, K. R. Brown, and I. L. Chuang. “Experimental investigation of planar ion traps”. *Phys. Rev. A - At. Mol. Opt. Phys.* **73** (2005). Cited on pages 10, 12.

## BIBLIOGRAPHY

- [Per03] A. Peres. *Finite precision measurement nullifies Euclid’s postulates*. 2003. arXiv: 0310035 [quant-ph]. Cited on page 221.
- [Pin21] J. M. Pino, J. M. Dreiling, C. Figgatt, J. P. Gaebler, S. A. Moses, M. S. Allman, C. H. Baldwin, M. Foss-Feig, D. Hayes, K. Mayer, C. Ryan-Anderson, and B. Neyenhuis. “Demonstration of the trapped-ion quantum CCD computer architecture”. *Nature* **592**, 209–213 (2021). Cited on pages 2, 10, 12.
- [Pir10] 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). Cited on page 228.
- [Pir14] S. Pironio, L. I. Masanes, A. Leverrier, and A. Acín. “Security of device-independent quantum key distribution in the bounded-quantum-storage model”. *Phys. Rev. X* **3** (2014). Cited on page 198.
- [Ple99] M. B. Plenio, S. F. Huelga, A. Beige, and P. L. Knight. “Cavity-loss-induced generation of entangled atoms”. *Phys. Rev. A* **59**, 2468–2475 (1999). Cited on page 288.
- [Pog21] I. Pogorelov, T. Feldker, C. D. Marciniak, L. Postler, G. Jacob, O. Kriegelsteiner, V. Podlesnic, M. Meth, V. Negnevitsky, M. Stadler, B. Höfer, C. Wächter, K. Lakhmanskiy, R. Blatt, P. Schindler, and T. Monz. *A compact ion-trap quantum computing demonstrator*. 2021. arXiv: 2101.11390. Cited on pages 149, 287.
- [Poh15] H. S. Poh, S. K. Joshi, A. Ceré, A. Cabello, and C. Kurtsiefer. “Approaching Tsirelson’s bound in a photon pair experiment”. *Phys. Rev. Lett.* **115** (2015). Cited on page 228.
- [Pop94] S. Popescu and D. Rohrlich. “Quantum nonlocality as an axiom”. *Found. Phys.* **24**, 379–385 (1994). Cited on page 199.
- [Poy96] J. F. Poyatos, J. I. Cirac, and P. Zoller. “Quantum reservoir engineering with laser cooled trapped ions”. *Phys. Rev. Lett.* **77**, 4728–4731 (1996). Cited on page 288.
- [Qia18] X. Qiang, X. Zhou, J. Wang, C. M. Wilkes, T. Loke, S. O’Gara, L. Kling, G. D. Marshall, R. Santagati, T. C. Ralph, J. B. Wang, J. L. O’Brien, M. G. Thompson, and J. C. Matthews. “Large-scale silicon quantum photonics implementing arbitrary two-qubit processing”. *Nat. Photonics* **12**, 534–539 (2018). Cited on page 3.
- [Rag20] S. Ragg. “Design and fabrication of a segmented ion trap optimized for cavity QED in the strong-coupling regime”. PhD thesis. ETH Zurich, 2020. Cited on page 337.



- [Rai92] M. G. Raizen, J. M. Gilligan, J. C. Bergquist, W. M. Itano, and D. J. Wineland. “Ionic crystals in a linear Paul trap”. *Phys. Rev. A* **45**, 6493–6501 (1992). Cited on page 8.
- [Ram14] R. Ramalingam, D. Boguhn, H. Fillinger, S. I. Schlachter, and M. Süßer. “Study of robust thin film PT-1000 temperature sensors for cryogenic process control applications”. In: *AIP Conf. Proc.* Volume 1573. 1. American Institute of Physics Inc., 2014, pages 126–131. ISBN: 9780735412019. Cited on page 182.
- [Ram50] N. F. Ramsey. “A molecular beam resonance method with separated oscillating fields”. *Phys. Rev.* **78**, 695–699 (1950). Cited on page 161.
- [Ree10] G. T. Reed, G. Mashanovich, F. Y. Gardes, and D. J. Thomson. “Silicon optical modulators”. *Nat. Photonics* **4**, 518–526 (2010). Cited on page 336.
- [Rei12] F. Reiter and A. S. Sørensen. “Effective operator formalism for open quantum systems”. *Phys. Rev. A - At. Mol. Opt. Phys.* **85**, 032111 (2012). Cited on page 288.
- [Rei16] F. Reiter, D. Reeb, and A. S. Sørensen. “Scalable Dissipative Preparation of Many-Body Entanglement”. *Phys. Rev. Lett.* **117**, 040501 (2016). Cited on page 288.
- [Rei17] F. Reiter, A. S. Sørensen, P. Zoller, and C. A. Muschik. “Dissipative quantum error correction and application to quantum sensing with trapped ions”. *Nat. Commun.* **8**, 1–11 (2017). Cited on page 288.
- [Reu17] K. Reuer. *Radio Frequency Voltage Sampling at Cryogenic Temperatures*. Technical report. Project report, ETH Zurich, 2017. Cited on page 83.
- [Rev20] M. C. Revelle. “Phoenix and Peregrine Ion Traps”. *arXiv* (2020). Cited on page 13.
- [Rie03] F. Riehle. *Frequency Standards*. Wiley, 2003, pages 1–526. ISBN: 9783527402304. Cited on pages 162, 166, 280.
- [Ron15] X. Rong, J. Geng, F. Shi, Y. Liu, K. Xu, W. Ma, F. Kong, Z. Jiang, Y. Wu, and J. Du. “Experimental fault-tolerant universal quantum gates with solid-state spins under ambient conditions”. *Nat. Commun.* **6**, 1–7 (2015). Cited on page 3.
- [Roo00] C. S. Roos, D. Leibfried, A. Mundt, F. Schmidt-Kaler, J. Eschner, and R. Blatt. “Experimental Demonstration of Ground State Laser Cooling with Electromagnetically Induced Transparency”. *Phys. Rev. Lett.* **85**, 5547–5550 (2000). Cited on page 49.
- [Roo08] C. F. Roos, T. Monz, K. Kim, M. Riebe, H. Häffner, D. F. James, and R. Blatt. “Nonlinear coupling of continuous variables at the single quantum level”. *Phys. Rev. A - At. Mol. Opt. Phys.* **77**, 040302 (2008). Cited on page 47.

## BIBLIOGRAPHY

- [Roo99] C. Roos, T. Zeiger, H. Rohde, H. C. Nägerl, J. Eschner, D. Leibfried, F. Schmidt-Kaler, and R. Blatt. “Quantum state engineering on an optical transition and decoherence in a paul trap”. *Phys. Rev. Lett.* **83**, 4713–4716 (1999). Cited on page 1.
- [Rou13] B. Rousseaux, S. Guérin, and N. V. Vitanov. “Arbitrary qudit gates by adiabatic passage”. *Phys. Rev. A - At. Mol. Opt. Phys.* **87**, 032328 (2013). Cited on page 234.
- [Rub06] E. Rubiola. *Tutorial on the double balanced mixer*. 2006. arXiv: 0608211 [physics]. Cited on page 117.
- [Rus16] T. Ruster, C. T. Schmiegelow, H. Kaufmann, C. Warschburger, F. Schmidt-Kaler, and U. G. Poschinger. “A long-lived Zeeman trapped-ion qubit”. *Appl. Phys. B Lasers Opt.* **122** (2016). Cited on pages 90, 169.
- [Sac00] C. A. Sackett, D. Kielpinski, B. E. King, C. Langer, V. Meyer, C. J. Myatt, M. Rowe, Q. A. Turchette, W. M. Itano, D. J. Wineland, and C. Monroe. “Experimental entanglement of four particles”. *Nature* **404**, 256–259 (2000). Cited on pages 1, 2.
- [San14] B. Sanguinetti, A. Martin, H. Zbinden, and N. Gisin. “Quantum random number generation on a mobile phone”. *Phys. Rev. X* **4**, 031056 (2014). Cited on page 229.
- [San18] S. Saner. *Diode injection stability Comparison between a Fabry-Perot diode and an anti-reflective coated diode*. Technical report. Project report, ETH Zurich, 2018. Cited on page 119.
- [Sca19] V. Scarani. *Bell nonlocality*. Oxford University Press, 2019. ISBN: 9780198788416. Cited on pages 201, 228.
- [Sch08] P. Schindler. *Frequency synthesis and pulse shaping for quantum information processing with trapped ions (MSc thesis)*. 2008. Cited on page 29.
- [Sch13] P. Schindler, D. Nigg, T. Monz, J. T. Barreiro, E. Martinez, S. X. Wang, S. Quint, M. F. Brandl, V. Nebendahl, C. F. Roos, M. Chwalla, M. Hennrich, and R. Blatt. “A quantum information processor with trapped ions”. *New J. Phys.* **15**, 123012 (2013). Cited on page 272.
- [Sch15] P. Schindler, D. J. Gorman, N. Daniilidis, and H. Häffner. “Polarization of electric-field noise near metallic surfaces”. *Phys. Rev. A* **92**, 013414 (2015). Cited on page 183.
- [Sed18a] J. A. Sedlacek, J. Stuart, W. Loh, R. McConnell, C. D. Bruzewicz, J. M. Sage, and J. Chiaverini. “Method for determination of technical noise contributions to ion motional heating”. *J. Appl. Phys.* **124**, 214904 (2018). Cited on page 12.

- [Sed18b] J. A. Sedlacek, J. Stuart, D. H. Slichter, C. D. Bruzewicz, R. McConnell, J. M. Sage, and J. Chiaverini. “Evidence for multiple mechanisms underlying surface electric-field noise in ion traps”. *Phys. Rev. A* **98**, 063430 (2018). Cited on pages 12, 181, 183.
- [Sel83] S. A. Self. “Focusing of spherical Gaussian beams”. *Appl. Opt.* **22**, 658 (1983). Cited on page 98.
- [Sen15] C. Senko, P. Richerme, J. Smith, A. Lee, I. Cohen, A. Retzker, and C. Monroe. “Realization of a quantum integer-spin chain with controllable interactions”. *Phys. Rev. X* **5**, 021026 (2015). Cited on page 334.
- [Sep12] M. Sepiol. *Frequency stabilization of a 729 nm diode laser to an external high finesse reference cavity (MSc thesis)*. 2012. Cited on page 113.
- [Sep19] 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”. *Phys. Rev. Lett.* **123**, 110503 (2019). Cited on page 162.
- [Set21] W. J. Setzer, M. Ivory, O. Slobodyan, J. W. Van Der Wall, L. P. Parazzoli, D. Stick, M. Gehl, M. Blain, R. R. Kay, and H. McGuinness. “Fluorescence Detection of a Trapped Ion with a Monolithically Integrated Single-Photon-Counting Avalanche Diode”. (2021). Cited on page 336.
- [Sha13] S. Shankar, M. Hatridge, Z. Leghtas, K. M. Sliwa, A. Narla, U. Vool, S. M. Girvin, L. Frunzio, M. Mirrahimi, and M. H. Devoret. “Autonomously stabilized entanglement between two superconducting quantum bits”. *Nature* **504**, 419–422 (2013). Cited on page 288.
- [She16] S. Sheldon, E. Magesan, J. M. Chow, and J. M. Gambetta. “Procedure for systematically tuning up cross-talk in the cross-resonance gate”. *Phys. Rev. A* **93**, 060302 (2016). Cited on page 3.
- [Sho94] P. Shor. “Algorithms for quantum computation: discrete logarithms and factoring”. In: *Proc. 35th Annu. Symp. Found. Comput. Sci.* IEEE Comput. Soc. Press, 1994, pages 124–134. ISBN: 0-8186-6580-7. Cited on page 1.
- [Sho95] P. W. Shor. “Scheme for reducing decoherence in quantum computer memory”. *Phys. Rev. A* **52**, R2493 (1995). Cited on page 1.
- [Sim13] S. H. Simon. *The Oxford Solid State Basics*. Oxford University Press, 2013. Cited on page 87.
- [Sim97] D. R. Simon. “On the power of quantum computation”. *SIAM J. Comput.* **26**, 1474–1483 (1997). Cited on page 1.
- [Siv11] J. D. Siversns, L. R. Simkins, S. Weidt, and W. K. Hensinger. “On the application of radio frequency voltages to ion traps via helical resonators”. *Appl. Phys. B Lasers Opt.* **107**, 921–934 (2011). Cited on page 82.

# BIBLIOGRAPHY

- [Sny83] A. W. Snyder and J. D. Love. *Optical Waveguide Theory*. Boston, MA: Springer US, 1983, pages 3–5. Cited on page 105.
- [Sør00] A. Sørensen and K. Mølmer. “Entanglement and quantum computation with ions in thermal motion”. *Phys. Rev. A - At. Mol. Opt. Phys.* **62**, 11 (2000). Cited on page 37.
- [Sør99] A. Sørensen and K. Mølmer. “Quantum computation with ions in thermal motion”. *Phys. Rev. Lett.* **82**, 1971–1974 (1999). Cited on page 37.
- [Spe05] R. W. Spekkens. *Contextuality for preparations, transformations, and unsharp measurements*. Technical report 5. 2005. arXiv: 0406166 [quant-ph]. Cited on page 223.
- [Spe14] R. W. Spekkens. *The Status of Determinism in Proofs of the Impossibility of a Noncontextual Model of Quantum Theory*. Technical report 11. 2014, pages 1125–1155. arXiv: 1312.3667. Cited on page 227.
- [Sri21] 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 two trapped ion qubits*. 2021. arXiv: 2102.12533. Cited on pages 2, 12.
- [Sta20] M. Stadler, V. Negnevitsky, U. Altunkaya, M. Malinowski, K. Mehta, C. Zhang, T.-L. Nguyen, C. Oenal, J. Home, M. Stadler, V. Negnevitsky, U. Altunkaya, M. Malinowski, K. Mehta, C. Zhang, T.-L. Nguyen, C. Oenal, and J. Home. *Multi-tone RF generation for trapped-ion control with low-latency feedback*. 2020. Cited on page 124.
- [Ste03] A. M. Steane. “Extracting Entropy from Quantum Computers”. *Ann. Henri Poincaré* **4**, 799–809 (2003). Cited on page 287.
- [Ste04] A. M. Steane. “How to build a 300 bit, 1 Giga-operation quantum computer”. *Quantum Inf. Comput.* **7**, 171–183 (2004). Cited on page 337.
- [Ste07] D. A. Steck. *Quantum and Atom Optics*. Technical report. 2007. Cited on pages 7, 15, 30.
- [Ste19] 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”. *Phys. Rev. Lett.* **124** (2019). Cited on page 106.
- [Ste86] S. Stenholm. “The semiclassical theory of laser cooling”. *Rev. Mod. Phys.* **58**, 699–739 (1986). Cited on pages 48, 49.
- [Ste96] A. M. Steane. “Error correcting codes in quantum theory”. *Phys. Rev. Lett.* **77**, 793–797 (1996). Cited on page 1.

- [Ste97] A. Steane. “The ion trap quantum information processor”. *Appl. Phys. B Lasers Opt.* **64**, 623–642 (1997). Cited on page 29.
- [Ste99] A. M. Steane. “Efficient fault-tolerant quantum computing”. *Nature* **399**, 124–126 (1999). Cited on page 1.
- [Sto03] A. I. Stognij, N. N. Novitskii, S. D. Tushina, and S. V. Kalinnikov. “Preparation of ultrathin gold films by oxygen-ion sputtering and their optical properties”. *Tech. Phys.* **48**, 745–748 (2003). Cited on page 336.
- [Stu18] J. Stuart, R. Panock, C. D. Bruzewicz, J. A. Sedlacek, R. McConnell, I. L. Chuang, J. M. Sage, and J. Chiaverini. “Chip-integrated voltage sources for control of trapped ions”. *Phys. Rev. Appl.* **11** (2018). Cited on pages 2, 13, 55.
- [Suc13] M. Suchara, A. Faruque, C.-Y. Lai, G. Paz, F. T. Chong, and J. Kubiatowicz. “Comparing the Overhead of Topological and Concatenated Quantum Error Correction”. *2013 IEEE 31st Int. Conf. Comput. Des. ICCD 2013* 419–426 (2013). Cited on page 3.
- [Sun19] Y. Sung, F. Beaudoin, L. M. Norris, F. Yan, D. K. Kim, J. Y. Qiu, U. von Lüpke, J. L. Yoder, T. P. Orlando, S. Gustavsson, L. Viola, and W. D. Oliver. “Non-Gaussian noise spectroscopy with a superconducting qubit sensor”. *Nat. Commun.* **10**, 1–8 (2019). Cited on page 272.
- [Tal15] I. Talukdar, D. J. Gorman, N. Daniilidis, P. Schindler, S. Ebadi, H. Kaufmann, T. Zhang, and H. Häffner. “Implications of surface noise for the motional coherence of trapped ions”. *Phys. Rev. A* **93** (2015). Cited on page 266.
- [Tan15] T. R. Tan, J. P. Gaebler, Y. Lin, Y. Wan, R. Bowler, D. Leibfried, and D. J. Wineland. “Multi-element logic gates for trapped-ion qubits”. *Nature* **528**, 380–383 (2015). Cited on page 2.
- [Tan17] T. R. Tan, Y. Wan, S. Erickson, P. Bierhorst, D. Kienzler, S. Glancy, E. Knill, D. Leibfried, and D. J. Wineland. “Chained Bell Inequality Experiment with High-Efficiency Measurements”. *Phys. Rev. Lett.* **118**, 130403 (2017). Cited on page 256.
- [Tel21] 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”. PhD thesis. 2021. arXiv: 2103.13846. Cited on pages 77, 184.
- [Tey10] F. Teyssandier and D. Prêle. *Commercially Available Capacitors at Cryogenic Temperatures*. Technical report. Ninth International Workshop on Low Temperature Electronics - WOLTE9, Jun 2010, Guarujá, Brazil., 2010. Cited on page 82.

## BIBLIOGRAPHY

- [Tho12] J. D. Thompson, T. G. Tiecke, A. S. Zibrov, V. Vuletić, and M. D. Lukin. “Coherence and Raman sideband cooling of a single atom in an optical tweezer”. *Phys. Rev. Lett.* **110** (2012). Cited on page 142.
- [Tic14] F. Ticozzi and L. Viola. “Steady-state entanglement by engineered quasi-local Markovian dissipation”. *Quantum Inf. Comput.* **14**, 265–294 (2014). Cited on page 288.
- [Tob16] W. G. Tobias, J. S. Rosenberg, N. R. Hutzler, and K. K. Ni. “A low-temperature external cavity diode laser for broad wavelength tuning”. *Rev. Sci. Instrum.* **87**, 113104 (2016). Cited on page 108.
- [Tod20] S. L. Todaro, V. B. Verma, K. C. McCormick, D. T. C. Allcock, R. P. Mirin, D. J. Wineland, S. W. Nam, A. C. Wilson, D. Leibfried, and D. H. Slichter. “State Readout of a Trapped Ion Qubit Using a Trap-Integrated Superconducting Photon Detector”. *Phys. Rev. Lett.* **126** (2020). Cited on pages 2, 134, 135, 336.
- [Tót14] G. Tóth and I. Apellaniz. *Quantum metrology from a quantum information science perspective*. 2014. arXiv: 1405.4878. Cited on page 35.
- [Tur00] Q. A. Turchette, C. J. Myatt, B. E. King, C. A. Sackett, D. Kielpinski, W. M. Itano, C. Monroe, and D. J. Wineland. “Decoherence and decay of motional quantum states of a trapped atom coupled to engineered reservoirs”. *Phys. Rev. A - At. Mol. Opt. Phys.* **62**, 053807–053801 (2000). Cited on pages 2, 24.
- [Tzi19] I. Tzitrin, J. E. Bourassa, N. C. Menicucci, and K. K. Sabapathy. “Progress towards practical qubit computation using approximate Gottesman-Kitaev-Preskill codes”. *Phys. Rev. A* **101** (2019). Cited on page 3.
- [Um13] M. Um, X. Zhang, J. Zhang, Y. Wang, S. Yangchao, D. L. Deng, L. M. Duan, and K. Kim. “Experimental certification of random numbers via quantum contextuality”. *Sci. Rep.* **3**, 012335 (2013). Cited on pages 228, 252, 253.
- [Um20] M. Um, Q. Zhao, J. Zhang, P. Wang, Y. Wang, M. Qiao, H. Zhou, X. Ma, and K. Kim. “Randomness Expansion Secured by Quantum Contextuality”. *Phys. Rev. Appl.* **10**, 34077 (2020). Cited on page 228.
- [Van10] A. P. VanDevender, Y. Colombe, J. Amini, D. Leibfried, and D. J. Wineland. “Efficient Fiber Optic Detection of Trapped Ion Fluorescence”. *Phys. Rev. Lett.* **105** (2010). Cited on page 134.
- [Van16] A. Van Rynbach, P. Maunz, and J. Kim. “An integrated mirror and surface ion trap with a tunable trap location”. *Appl. Phys. Lett.* **109**, 221108 (2016). Cited on page 270.
- [vBur20] V. von Burg, G. H. Low, T. Häner, D. S. Steiger, M. Reiher, M. Roetteler, and M. Troyer. *Quantum computing enhanced computational catalysis*. 2020. arXiv: 2007.14460. Cited on page 3.

- [vdHar03] H. W. van der Hart. “Two-photon double ionization of helium”. In: *Proceedings*. Volume 697. 1. AIP Publishing, 2003, pages 27–36. Cited on page 149.
- [Ver09] F. Verstraete, M. M. Wolf, and J. Ignacio Cirac. “Quantum computation and quantum-state engineering driven by dissipation”. *Nat. Phys.* **5**, 633–636 (2009). Cited on pages 3, 287, 288.
- [Vit13] G. Vittorini. “Stability of ion chains in a cryogenic surface-electrode ion trap”. PhD thesis. Georgia Tech, 2013. Cited on page 258.
- [vMou21] M. W. van Mourik, P. Hrmo, L. Gerster, B. Wilhelm, R. Blatt, P. Schindler, and T. Monz. *RF-induced heating dynamics of non-crystallized trapped ions*. 2021. arXiv: 2104.10623. Cited on page 145.
- [Waj16] M. Wajs, S. Y. Lee, P. Kurzyński, and D. Kaszlikowski. “State-recycling method for testing quantum contextuality”. *Phys. Rev. A* **93**, 6 (2016). Cited on pages 228, 239.
- [Wan10] C. Wang and J. G. Hartnett. “A vibration free cryostat using pulse tube cryocooler”. *Cryogenics (Guildf)*. **50**, 336–341 (2010). Cited on page 56.
- [Wan11] S. X. Wang, G. H. Low, N. S. Lachenmyer, Y. Ge, P. F. Herskind, and I. L. Chuang. “Laser-induced charging of microfabricated ion traps”. *J. Appl. Phys.* **110** (2011). Cited on pages 12, 270.
- [Wes19] G. N. West, W. Loh, D. Kharas, C. Sorace-Agaskar, K. K. Mehta, J. Sage, J. Chiaverini, and R. J. Ram. “Low-loss integrated photonics for the blue and ultraviolet regime”. *APL Photonics* **4**, 26101 (2019). Cited on page 335.
- [Win14] A. Winter. *What does an experimental test of quantum contextuality prove or disprove?* Technical report 42. 2014. arXiv: 1408.0945. Cited on pages 222, 228.
- [Win97] 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”. *J. Res. Natl. Inst. Stand. Technol.* **103**, 259–328 (1997). Cited on pages 1, 7, 14, 46.
- [Win98] D. Wineland, C. Monroe, W. Itano, B. King, D. Leibfried, D. Meekhof, C. Myatt, and C. Wood. “Experimental Primer on the Trapped Ion Quantum Computer”. *Fortschritte der Phys.* **46**, 363–390 (1998). Cited on page 265.
- [Woo11] C. J. Wood, J. D. Biamonte, and D. G. Cory. “Tensor networks and graphical calculus for open quantum systems”. *Quantum Inf. Comput.* **15**, 759–811 (2011). Cited on page 293.
- [Woo82] W. K. Wootters and W. H. Zurek. “A single quantum cannot be cloned”. *Nature* **299**, 802–803 (1982). Cited on page 1.

## BIBLIOGRAPHY

- [Wri13] K. Wright, J. M. Amini, D. L. Faircloth, C. Volin, S. Charles Doret, H. Hayden, C. S. Pai, D. W. Landgren, D. Denison, T. Killian, R. E. Slusher, and A. W. Harter. “Reliable transport through a microfabricated X-junction surface-electrode ion trap”. *New J. Phys.* **15**, 033004 (2013). Cited on page 10.
- [Xie17] X. Xie, R. Bouchand, D. Nicolodi, M. Lours, C. Alexandre, and Y. L. Coq. “Phase noise characterization of sub-hertz linewidth lasers via digital cross correlation”. *arXiv* (2017). Cited on page 113.
- [Yog19] M. Yoganathan and C. Cade. “The one clean qubit model without entanglement is classically simulable”. (2019). Cited on page 4.
- [Yu12] S. Yu and C. H. Oh. “State-independent proof of Kochen-Specker theorem with 13 rays”. *Phys. Rev. Lett.* **108**, 1–5 (2012). Cited on page 216.
- [Zha13] X. Zhang, M. Um, J. Zhang, S. An, Y. Wang, D. L. Deng, C. Shen, L. M. Duan, and K. Kim. “State-independent experimental test of quantum contextuality with a single trapped ion”. *Phys. Rev. Lett.* **110**, 1–5 (2013). Cited on page 239.
- [Zha17] Z. Zhao, R. Pisarczyk, J. Thompson, M. Gu, V. Vedral, and J. F. Fitzsimons. “Geometry of quantum correlations in space-time”. *Phys. Rev. A* **98** (2017). Cited on page 193.
- [Zha21] C. Zhang. “Multi-ion quantum logic on a planar ion-trap with integrated photonics”. PhD thesis. ETH Zurich, 2021. Cited on pages 23, 76, 83, 84, 142, 147, 153, 171, 314, 319, 333.
- [Zho17] P. Zhou. *Motional State Control and Ablation Loading in a Cryogenic Surface-Electrode Trap (MSc thesis)*. 2017. Cited on page 130.
- [Zu12] C. Zu, Y. X. Wang, D. L. Deng, X. Y. Chang, K. Liu, P. Y. Hou, H. X. Yang, and L. M. Duan. “State-independent experimental test of Quantum contextuality in an indivisible system”. *Phys. Rev. Lett.* **109**, 1–5 (2012). Cited on page 239.
- [Zur20] M. Zurel, C. Okay, and R. Raussendorf. “Hidden Variable Model for Universal Quantum Computation with Magic States on Qubits”. *Phys. Rev. Lett.* **125**, 260404 (2020). Cited on page 4.