ABSTRACT

Title of dissertation: Simulating many-body quantum spin models with trapped ions

Antonis Kyprianidis
Doctor of Philosophy, 2021

Dissertation directed by: Professor Christopher Monroe
Department of Physics

Richard Feynman in 1981 suggested using a quantum machine to simulate quantum mechanics. Peter Shor in 1994 showed that a quantum computer could factor numbers much more efficiently than a conventional one. Since then, the explosion of the quantum information field is attesting to how motivation and funding work miracles. Research labs in the field are multiplying, commercial companies manufacturing prototypes are proliferating, undergraduate Physics curricula incorporate more than one courses in aspects of quantum information, quantum advantage over classical computers has been claimed, and the United States and European Union will be spending more than $\$10^9$ each in quantum information over the next few years. Naturally, this expansion has led to diversification of the devices being developed. The quantum information systems that cannot simulate an arbitrary evolution, but are specialized in a specific set of Hamiltonians, are called quantum *simulators*. They enjoy the luxury of being able to surpass computational abilities of classical computers *right now*, at the expense of only doing so for a narrow type

of problem. Among those systems, ions trapped in vacuum by electric fields and manipulated with light have proved to be a leading platform in emulating quantum magnetism models. In this thesis I present trapped-ion experiments realizing a prethermal discrete time crystal. This exotic phase occurs in non-equilibrium matter subject to an external periodic drive. Normally, the ensuing Floquet heating maximizes the system entropy, leaving us with a trivial, infinite-temperature state. However, we are able to parametrically slow down this heating by tuning the drive frequency. During the time window of slow thermalization, we define an order parameter and observe two different regimes, based on whether it spontaneously breaks the discrete time translation symmetry of the drive or it preserves it. Furthermore, I demonstrate a simple model of electric field noise classically heating an ion in an anharmonic confining potential. As ion traps shrink, this kind of noise may become more significant. And finally, I discuss a handful of error sources. As quantum simulation experiments progress to more qubits and complicated sequences, accounting for system imperfections is becoming an integral part of the process.

Simulating many-body quantum spin models with trapped ions

by

Antonis Kyprianidis

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2021

Advisory Committee:
Professor Chris Monroe, Chair/Advisor
Professor Norbert Linke
Professor Alexey Gorshkov
Professor Ronald Walsworth
Professor Mohammad Hafezi

© Copyright by
Antonis Kyprianidis
2021

Dedication

To Cara

Acknowledgments

First and foremost, I thank my advisor, Chris Monroe, for welcoming me in the group with no prior experimental experience. He encouraged me to focus on what matters and always offered a broad perspective of how science has evolved, and the forces that steer it. I am also grateful for witnessing in practice that our scientific research is not a zero-sum game; on the contrary, multiple people and groups can be working toward the same result while benefiting from each other's existence in the long run.

A special thanks goes to Will Morong for being an outstanding mentor, despite our relatively short temporal overlap. His scientific acuity and passion for mentorship have made a difference for me, while his patience, leadership, and thought-out arguments have been inspiring.

Jiehang, Aaron, Jake, thank you for guiding me through the first phase of my PhD and teaching me how to solder, how to couple a laser beam to a fiber, and how to align a mirror. You tried to teach me much more advanced things than those, too, but I'm not sure I understood at the time; I think it still benefited me down the road though! Guido, grazie for all the guidance and availability in the second half of my degree—QSim would be different without you!

My labmates Patrick and Kate: many days I would see them more than I saw my wife. A good team. *The show must go on* no matter what, but this was way easier with collaborators showing integrity, friendliness, and compassion. Bonding over shared frustrations is real!

I am also indebted to the rest of the ion trapper ecosystem in UMD. Be it for a quick word of support or patiently passing down knowledge, I have come to appreciate the collective interest and team spirit of the UMD team. To name just a few: Norbert L., Qudsia Q., Joe B., Kristi B., Marty L., Paul H., Crystal N., David W. C., Harvey K., Caroline F., Shantanu D., Kale J., Jonathan M., Volkan I., Ken W., Clay C., Alaina G., Lei F., Cinthia H.A., Yingyue Z., Nhung N., Arinjoy D., Debo V., Jameson O'R., Sagnik S.—many thanks. Marko C., thank you for always taking the time to share your knowledge. Additionally, I need to thank theory collaborators and most notably Norm Yao, Francisco Machado, Dominic Else, Chris Jarzynski, and Wade Hodson, for giving us plenty of good ideas to work with, and patiently explaining them to us first.

As far as our AMO research goes, it cannot be done in a cave. All employees maintaining campus a pleasant and productive environment are the often under-appreciated base. While not integral to research, the Stamp ice cream shop employees have also literally sweetened up many hectic days for me, even trusting I will come back and pay that day I had forgotten my wallet.

Words can hardly describe how grateful I am to my family for always being supportive and one Skype call away: my parents Kornilio and Evi, my sister and brother-in-law Sofianna and Nikola. Among so many other virtues, their dignity, perseverance, and stoicism set the example for me. Also, my in-laws Lori, Will, Julia, James, Nate, and Avi, for making me feel like part of their family with their warmth and generosity, long before I actually join it.

Wen Lin and Daiwei, thank you for being friends through good and bad times,

for sharing and listening. Of course my Greek friends too and the long phone calls that were often reminding me of my pre-grad school life: Giorgo, Mano, Thoma, Lambro, many thanks.

When research is not going where you want, or as fast as you would like, it's vital to maintain other sources of happiness providing that valuable resource. Oddly for a physicist, I have mined so much fun from the salsa dancing community of the DC area that I feel it truly deserves a thanks—especially my mentor and instructor Barbara Bernstein.

Lastly, I need to thank Cara, the woman I met on a campus bus stop at the beginning of my PhD. Six years, five miles, and one pandemic down the road, we became partners in marriage. She has been an inspiration and a support to an extent I did not think possible. For better, for worse, in sickness and in health, in lab-crunching weeks and holidays, she has been a true companion. Her positive attitude, multi-layered perception of reality, and organizational skills keep expanding my point of view. I owe her a lot, and this thesis is dedicated to her.

Table of Contents

List of Tables

List of Figures

Chapter 1: Introduction

In 1989, Wolfgang Paul would be awarded the Physics Nobel prize "for the development of the ion trap technique", the world would see the Berlin wall demolished, and I started my journey on Earth blissfully unaware of the social and technological developments that would mark my generation. Thirty years later, this PhD thesis is presenting aspects of quantum simulation using an ion trap, offering a perspective on the progress made and current challenges. One wonders what PhD theses in quantum information—if the field is still called so—will be about in another thirty years. Or perhaps sooner. If social and commercial interest has spurred Moore's law of exponential growth in conventional computing power[1], it remains to be seen what it will do for quantum information devices.

Within the field of quantum information is quantum simulation. Experimental quantum simulators are composed of a diverse landscape of devices [1] that mimic specific models (Hamiltonians) with some degree of parameter tunability. Underlying the simulations in our lab is the following observation: the physical description of laser beams interacting with trapped ions can be made similar to the description of a set of quantum spins interacting with one another, optionally at the presence

[1] More accurately, exponential growth in the number of components in a dense integrated circuit.

of magnetic fields.

A number of experiments has been made possible in this platform [2, 3], from adiabatically prepared ground states of spin Hamiltonians [4] and information spreading velocities in long-range interacting quantum systems [5], to demonstrations of lack of thermalization [6, 7] and discrete time crystals [8], with a maximum number of 53 qubits [9] for published research. This breadth reflects the ability to control and measure the spins, and the tuneability of the applied Hamiltonian.

In this chapter I will mention some key elements of our experiments to merely set the scene for the reader. For extended descriptions of the physical apparatus or the derivation of how our simulation Hamiltonian arises from the physical Hamiltonian, I point the reader to references included and graduate theses.

1.1 An ion, our qubit

The ion species we use is $^{171}\mathrm{Yb}^+$. Each ion encodes an effective spin-1/2 in its hyperfine levels $|\downarrow\rangle_z \equiv |F = 0, m_F = 0\rangle$ and $|\uparrow\rangle_z \equiv |F = 1, m_F = 0\rangle$ of the $^2S_{1/2}$ manifold. F is the quantum number for the total angular momentum of the ion and m_F is the projection on the quantization axis, set by a magnetic field of magnitude ≈ 5 Gauss. Being able to *encode* an effective spin-1/2 means that during our simulation, we maintain a two-level system, and all other levels have been somehow "eliminated" from the simulation picture. This elimination happens by choosing to always couple to our ion with laser light that may physically be using states other than $|\uparrow\rangle_z$ and $|\downarrow\rangle_z$, but it can be effectively described without needing

those, merely using Pauli spin matrices acting on $|\downarrow\rangle_z$ and $|\uparrow\rangle_z$.

The hyperfine splitting $\omega_{\mathrm{HF}} \approx 2\pi \cdot 12.64$ GHz between $|\downarrow\rangle_z$ and $|\uparrow\rangle_z$ is in the RF spectrum. Before each experimental sequence, the ions are Doppler-cooled with 369 nm light, initialized in the $|\downarrow\rangle_z$ state with optical pumping, and Raman sideband-cooled to the motional ground state [10]. After the main experimental sequence is finished (see section 1.3), the state of the ion/spin is projectively measured with 369 nm light. Since each ion's aforementioned levels are mapped to a single spin, I will often use the terms "ion" and "spin" interchangeably in this thesis.

1.1.1 State initialization, aka optical pumping

At the beginning of every experimental sequence, right after Doppler-cooling our ions, we initialize their internal state to ${}^2S_{1/2}|F = 0, m_F = 0\rangle$, i.e. the qubit $|\downarrow\rangle_z$ state. The complete procedure is described in [10] or S. Olmschenk's thesis [11]. To intuitively understand how we are able to efficiently and accurately initialize our state to $|\downarrow\rangle_z$, first let's see what states we are likely to have after Doppler cooling, and second, how the optical pumping mechanism takes all of them indiscriminately to the $|\downarrow\rangle_z$ state.

Right after Doppler cooling, the ion is most likely in its ground state, the ${}^2S_{1/2}$ manifold (the 4 states shown in Fig. 1.1). Other states are unlikely, since most decay fast to one of those S states via spontaneous emission.

Keeping in mind this manifold of possible starting states, let's see what happens when we shine the optical pumping light. This 369.5 nm light is resonant with

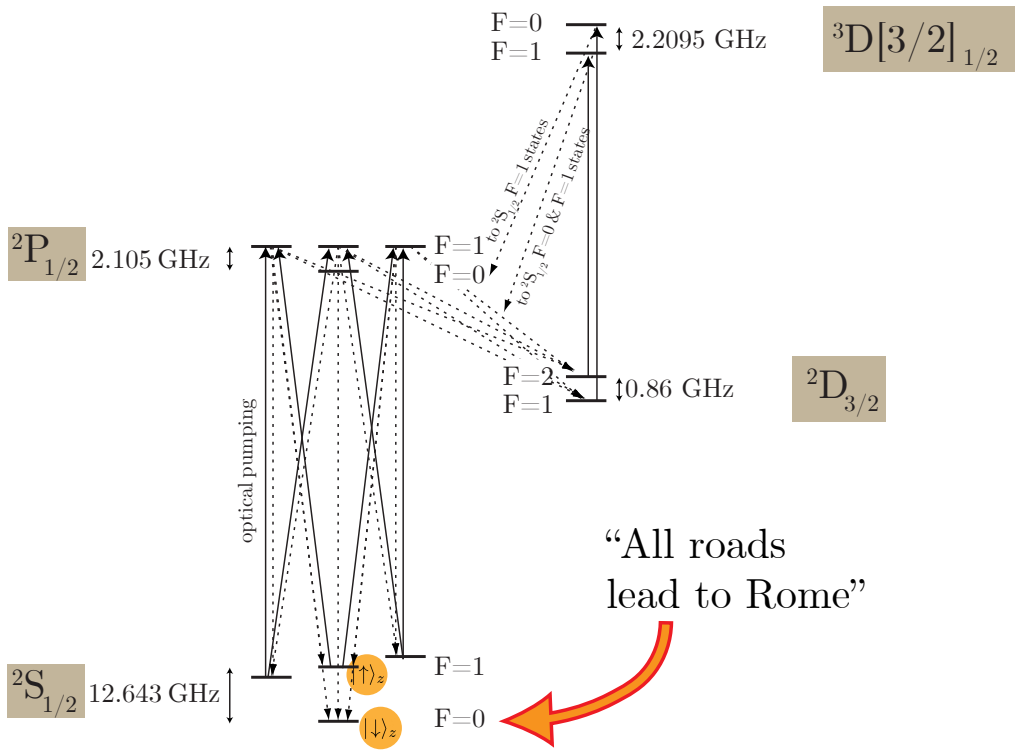

Figure 1.1: **State initialization with optical pumping.** All roads lead to Rome: no matter which state of the S manifold the ion starts from after Doppler cooling, the optical pumping light will eventually couple it to the $|0\rangle = |\downarrow\rangle_z$ state, after a probabilistic small number of photon absorptions and spontaneous decays. Solid lines indicate stimulated absorption from the optical pumping light, dotted lines show spontaneous decay. Adapted from [10].

the $|\downarrow\rangle_z \leftrightarrow P_{1/2}|F=1, m_F=1\rangle$, and has all three polarization components relative to the quantizing field. With respect to this light...

- ...the $|\downarrow\rangle_z$ state is 12.6 GHz-detuned from the $P_{1/2}|F=1\rangle$ manifold, and the transition $|\downarrow\rangle_z \rightarrow P_{1/2}|F=0\rangle$ is dipole-forbidden. So, the $|\downarrow\rangle_z$ state remains unchanged by the optical pumping light.

- ...the $|\uparrow\rangle_z$ state is forbidden from selection rules to couple to the same F and same m_F state $P_{1/2}|F=1, m_F=0\rangle$. Therefore it can couple to either $P_{1/2}|F=1, m_F=-1\rangle$ or $P_{1/2}|F=1, m_F=1\rangle$. From each of those, there is a 1/3 probability to decay to $|\downarrow\rangle_z$ and get trapped there (see point above). The rest 2/3 of the time these two P states decay either back to $|\uparrow\rangle_z$, in which case this point "loops" to the beginning, or to one of the Zeeman levels of the S manifold, in which case...

- ...the Zeeman states $S_{1/2}|F=1, m_F=-1\rangle$ and $S_{1/2}|F=1, m_F=1\rangle$ can each couple to the same Zeeman state of the $P_{1/2}$ manifold, or the $P_{1/2}|F=1, m_F=0\rangle$. From there, they can decay again with a probability of 1/3 to the deisred $|\downarrow\rangle_z$ state, or another state of the S manifold, in which case we loop to one of the previous points.

The conclusion is that the only *final* state from which the ion cannot escape during optical pumping is the desired, $|\downarrow\rangle_z$. Any other state it couples to will keep coupling to other states, with a small leakage probability at each decay taking it to $|\downarrow\rangle_z$, and staying there. So...one way or the other, all possible paths (roads) will eventually lead the ion to $|\downarrow\rangle_z$ (Rome)!

A caveat is that when the ion is at the desired $|\downarrow\rangle_z$ state, there is a very small probability that the optical pumping light will excite it to the $P_{1/2}|F=1\rangle$ manifold, from which it is detuned by $12.6 + 2.1 = 14.7$ GHz.

Finally, the last step needed for optical pumping is to account for the 0.5% probability that an ion in any state of the $P_{1/2}$ manifold will decay not to somewhere in the $S_{1/2}$ manifold, but to the $^2D_{3/2}$ state. The optical pumping light cannot couple $D_{3/2}$ to anywhere, so we need another laser to prevent the ion from getting stuck there. This laser is a 935 nm resonant to the $D_{3/2}|F=2\rangle \leftrightarrow D[3/2]_{1/2}|F=1\rangle$. We add sidebands at $2.2095 + 0.86 = 3.0695$ GHz to also couple $D_{3/2}|F=1\rangle$ to $D[3/2]_{1/2}|F=0\rangle$. From the $D[3/2]_{1/2}$ state, the ion decays back to the S manifold with spontaneous decay time of 37.7 ns [12].

1.1.2 State readout

After preparing the ion qubits in the $|\downarrow\rangle_z$ state, we use the Raman beams (see section later) to perform the main part of the simulation. At the end of it, we are interested in measuring their state. To this end, we illuminate all of the ion chain with 369.5 nm light, resonant with the $S_{1/2}|F=1, m_F=0\rangle \leftrightarrow P_{1/2}|F=0, m_F=0\rangle$ transition and containing all three polarizations $(\sigma_-, \sigma_+, \pi)$ with respect to the quantizing field. A nice overview of the procedure is given in [13] and [14], as well as [15] and [16] for a similar system in $^{40}Ca^+$ ions. The main idea (Fig. 1.2) is that if an ion is projected to the $|\downarrow\rangle_z$ state, the detection will not couple it to any state and therefore the ion will appear dark. If on the other hand the ion is projected to

the $|\uparrow\rangle_z$ state, it enters a cycling transition, constantly absorbing and re-emitting 369.5 nm photons. Using an objective lens system with NA = 0.4, we collect part of these photons, imaging the ion chain on the sensor of an EMCCD camera.

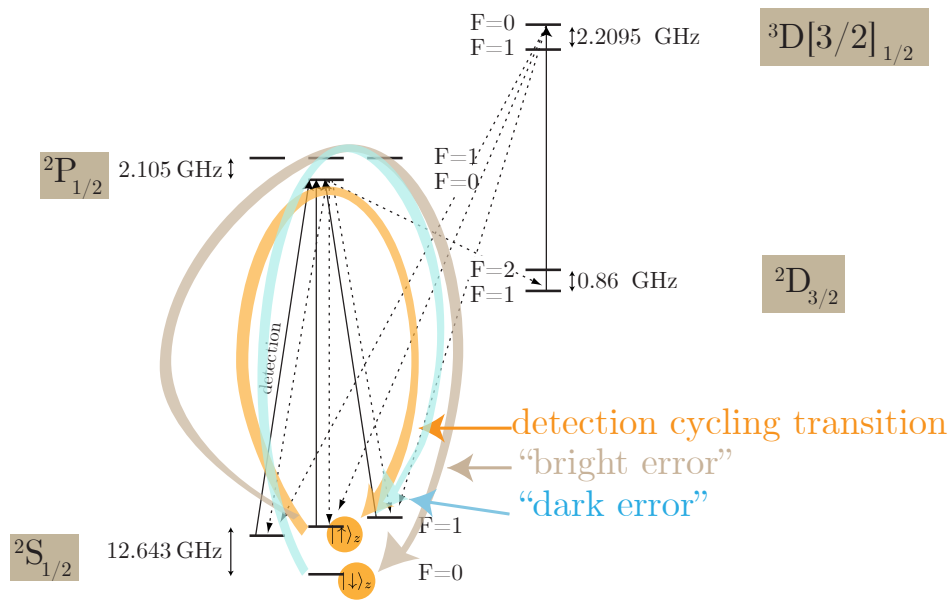

Figure 1.2: **State readout cycling transition and error modes.** Energy levels relevant to state readout. Similar to optical pumping, we need the 935 nm repump laser. If an ion is projected to $|\uparrow\rangle_z$ by the resonant detection beam, it enters the orange cycling transition: photon absorption→photon scattering→ and we collect these fluorescence photons with our imaging optics. There is a finite probability that during this cycle, an ion initially projected to $|\uparrow\rangle_z$ will off-resonantly couple to the $P_{1/2}|F = 1\rangle$ manifold, detuned by 2.1 GHz from the light frequency. From there, it may decay to the $|\downarrow\rangle_z$ dark state and stop fluorescing; a "bright error" has occurred. Additionally, an ion initially projected to the dark $|\downarrow\rangle_z$ state may couple to the same $P_{1/2}|F = 1\rangle$ manifold and decay to $|\uparrow\rangle_z$. This transition is less likely to occur since it is detuned by $12.6 + 2.1 = 14.7$ GHz from the light frequency. This is a "dark error". Adapted from [10].

There are a few things that can introduce errors at the readout stage. First, the detection light might off-resonantly couple the state the ion was initially projected at to the $|F = 1\rangle$ manifold of the $P_{1/2}$ state, and from there decay to a different qubit state. The process is illustrated in Fig. 1.2. This can happen any time during

the exposure (i.e., detection) time window of $0.4 - 0.65$ ms, and will appear as an ion with fluorescence between zero and maximum. Typical numbers for these probabilities are 1% for the dark error (erroneous flip $|\downarrow\rangle_z \rightarrow |\uparrow\rangle_z$) and 3% for the bright error (erroneous flip $|\uparrow\rangle_z \rightarrow |\downarrow\rangle_z$). Secondly, during the exposure time, 369 nm light directly from the laser may enter the imaging optics and fall into the region of the camera sensor where ion fluorescence would head. In this case, our state discrimination algorithm may no be able to tell the difference. Third, the camera sensor comprises of pixels, any of which might "fire" at any moment and produce a signal as if it had received a photon, even if it didn't. Reducing this noise is the primary reason we keep our camera sensor at a low temperature of $-60°$ C.

As an aside, when the ion is not properly pumped to the $|\downarrow\rangle_z$ state, a lobe will appear in the dark and bright count histograms, taken in the calibration stage. Fig. 1.3 illustrates a time when this happened. The cause for this bad optical pumping could be the optical pumping light itself, or the 935 nm repump light, including their sidebands. In Fig. 1.3, we see the difference between problematic and functional optical pumping, as demonstraded in the bright and dark count histograms. Note that the off-resonant optical pumping during the exposure (detection) time mentioned in the previous paragraph will not resut in a lobe appearing in the histograms, but rather in the flat extension of, say, the bright counts histogram into the dark counts region and vice versa.

For an overview of how we use the ion fluorescence to decide if a certain ion was at the $|\uparrow\rangle_z$ or $|\downarrow\rangle_z$ state, see Appendix B.

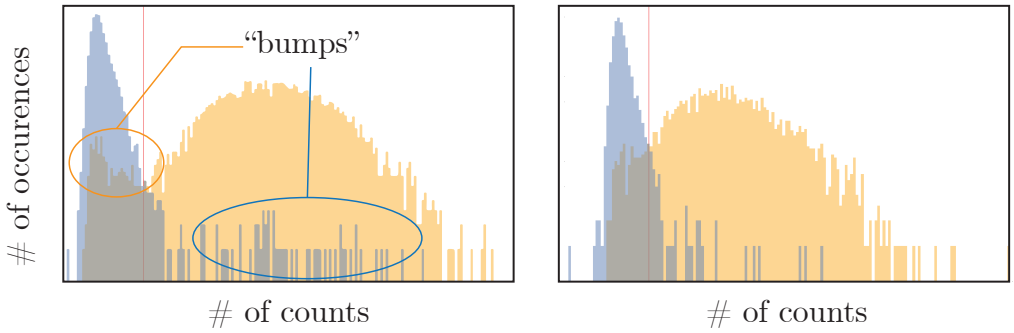

Figure 1.3: **The signature of bad optical pumping in the detection calibration histograms. Left:** The two bumps appearing in the bright (yellow) and the dark (blue) counts histogram indicate an issue with state preparation. For the dark counts, the ions undergo the sequence Doppler cooling → Optical pumping → Detection, while for the bright counts, they undergo Doppler cooling → Optical pumping → Raman sideband cooling → π rotation to $|\uparrow\rangle_z$ → Detection. **Right:** after optimizing the optical pumping issue, the lobes disappear.

1.1.2.1 Imaging system magnification

To estimate the overall magnification of our imaging system, we first measured the axial trap frequency using a single ion. Then, we loaded 20 ions and measured the minimum spacing in pixel units. We then calculated numerically the expected minimum ion spacing at that axial confinement, and divided that by the spacing measured in pixels. The conversion factor is

$$1 \text{ camera pixel} \leftrightarrow 0.23 \pm 0.02\ \mu m.$$

Knowing that the camera pixels are 16×16 μm in size, we estimate the magnification of our system to be roughly 70×.

1.1.3 Single-qubit manipulation: The Raman beams

In-between initializing and measuring our qubits is where the magic, aka quantum simulation, happens. The operations that the spins undergo during the simulation stage fall under two different categories: single-body terms σ_i^q and many-body terms $\sim \sigma_i^q \sigma_j^{q'}$, where σ_i^q is the q-th Pauli spin matrix acting on qubit i. I will refer to the many-body terms in the next section with more detail. As for the single-body terms, in simple language, they are rotations of the spin-1/2 vector around the Bloch sphere. As a reminder, a pure state $|\psi\rangle$ of a two-level system can always be written as

$$|\psi\rangle = \cos\frac{\Theta}{2}|\downarrow\rangle_z + e^{i\Phi}\sin\frac{\Theta}{2}|\uparrow\rangle_z, \qquad |\langle\psi|\psi\rangle|^2 = 1,$$

allowing us to define the vector $\{\Theta, \Phi\}$ in spherical coordinates.

The ability to perform these Bloch rotations boils down to the ability to manipulate the qubit states $|\downarrow\rangle_z$ and $|\uparrow\rangle_z$, creating arbitrary superpositions of them. We do that by using Raman transitions between the corresponding atomic states. To drive these transitions, two Raman beams ("Raman 1" and "Raman 2") at a 90° angle with respect to each other and 45° each relative to the ion chain converge at the ions' location. Their frequency difference is tuned to be about $f_A = 29.7$ MHz using Brimrose AOMs. They come from a pulsed laser (Coherent Paladin Compact 4W, air-cooled) with a repetition rate $f_{\text{rep}} \approx 120$ MHz, which means that each of those beams is an optical frequency comb. For an overview of frequency combs as used in our setup, see Appendix A.2. The frequency difference f_A ensures that

every optical comb tooth from Raman 1 interferes with an optical comb tooth from Raman 2 that is 105 teeth away in frequency and results in an RF frequency exactly equal to the hyperfine splitting between $|\downarrow\rangle_z$ and $|\uparrow\rangle_z$:

$$105 f_{\text{rep}} + f_A = \omega_{\text{HF}}.$$

The polarization of the Raman beams is horizontal (parallel to the table), as per

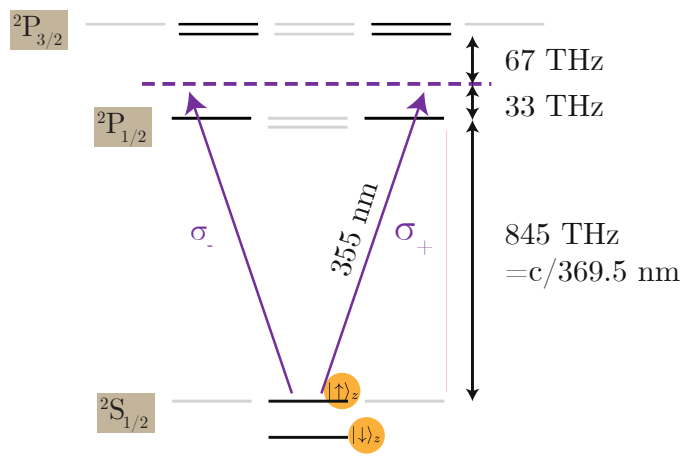

Figure 1.4: **Raman transitions and the relevant energy levels.** Each of the Raman beams is an equal superposition of σ_- and σ_+ polarizations, coupling the qubit levels to the $|m_F = \pm 1\rangle$ levels of the $P_{1/2}$ and the $P_{3/2}$ manifolds. These levels act as intermediate (adiabatically eliminated) states for the Raman transition to the other qubit level.

the "lin⊥lin" configuration. For reference, we can write their polarizations ϵ_i in the $(\sigma_-, \pi, \sigma_+)$ notation as

$$\epsilon_1 = \frac{1}{2}(1 - i, 0, 1 + i) \tag{1.1}$$

$$\epsilon_2 = \frac{1}{2}(1 + i, 0, 1 - i) \tag{1.2}$$

with $|\epsilon_i|^2 = 1$.

Table 1.1: Reference table of Raman beam parameters[1]

Carrier Rabi frequency Ω_0		880 kHz [2]	
Optical power in each Raman 1st order		790 mW	
Raman 1		Raman 2	
Horizontal waist	119 μm	Horizontal waist	102 μm
Vertical waist	9.8 μm	Vertical waist	11.9 μm

[1] As measured on October 2, 2020.
[2] At maximum power for the Raman 2 beam.

The Raman beams induce Rabi flopping between the qubit states at a Rabi frequency Ω_0 proportional to the geometric mean of their intensities I_1 and I_2 at the ion location:

$$\Omega_0 \sim \sqrt{I_1 I_2}$$

Given that the sum of their intensities is constant, the geometric mean is maximized when $I_1 = I_2 \equiv I$ and then

$$\Omega_0 \sim I$$

At the ions location, the Raman beams assume a gaussian profile with waists roughly 100 μm and 10 μ along the horizontal and vertical axes respectively. However, this profile changes periodically because of optics misalignment or degradation and we have to check it often. For reference, some related measurements are shown in Table 1.1.

The spatial dependence of the intensity because of the gaussian beam profile created two issues deemed worth addressing in the past few years, in order to improve our operation fidelities. First, beam pointing noise of the beams at the ion location

was translating into intensity noise and therefore noise in Ω_0[2]. Secondly, for large ion chains that span a length comparable to the beams' horizontal waist, the beam intensity sampled varies significantly and so does Ω_0. Therefore, if we apply the Raman beams for a duration that is a π-time around some axis for the central ion, the ions at the edge of the chain are rotated by less than π. We addressed the first issue by implementing a PI lock of the beam pointing and redesigning the Raman beam path to be telecentric. The second issue was addressed by replacing simple Bloch rotations with a composite pulse sequence of four rotations. The following two sections provide more details.

1.1.3.1 The Raman beam pointing lock

In order to stabilize the beam pointing at the ion location, we changed our Raman beam path and implemented a PI stabilization scheme. As shown in Fig. 1.5, the 355 nm light exits the Paladin laser roughly collimated, goes through a $\times 1.33$ telescope, and then gets slowly focused into the Raman AOM centers by a spherical +412 mm lens, to a focus of 160-200 μm radius. The lens performs a spatial Fourier transform of the wavefront from the object point (inside the laser) to the image at the AOMs. Therefore, all angular noise out of the laser is mapped to displacement noise at the AOMs.

This displacement noise is easily measured by sampling the beam and directing it to a quadrant photodiode (QPD)[3] at the effective AOM location. The QPD

[2]It also affects noise in the AC Stark shift, but as I'll show in Chapter 2, pointing noise is not the primary source of AC Stark shift noise in our system.

[3]QD50-0-SD from OSI Optoelectronics.

circuitry has two voltage outputs, one for the horizontal and one for the vertical axis. Each voltage output is proportional to the distance of the beam from the diode sensor center along that axis. The sensor has a dead area ("element gap") in its center at the shape of a cross, 0.2 mm wide. While this area receives most of the beam intensity, the part of the gaussian intensity profile outside this region is sufficient to provide a reliable input signal. The QPD output signal is amplified[4] and sent as the error signal to a PI lock that uses a half-inch mirror upstream as an actuator; see Fig. 1.5. A half-inch mirror was chosen, as opposed to a one-inch, due to its faster response to repositioning, owing to smaller size. Both the horizontal and the vertical knob of the mirror are controlled by piezoelectric actuators. The displacement noise at the sampling position is reduced by almost 100 when the lock is ON and tuned.

After the AOMs, the diverging beams are collimated by spherical +240 mm lenses, placed one focal length downstream from the AOMs. Up to this point, their profile is spherical; beyond these lenses, their horizontal and vertical axes are shaped separately to ensure they have a large horizontal waist at the ion location to address long ion chains, and a small vertical one so as not to waste light intensity in the vertical axis. Because of the high density of optical elements, it has not been possible to fully enclose the Raman beam path with beam tubes, which would reduce its susceptibility to air turbulence. Therefore, the 100-fold reduction of the pointing noise mentioned does not carry over to the ion location. Measurements indicated a reduction by a factor between 3 and 10 at the ion location when the lock

[4]Using the AD8421ARZ-ND instrumentation amplifier.

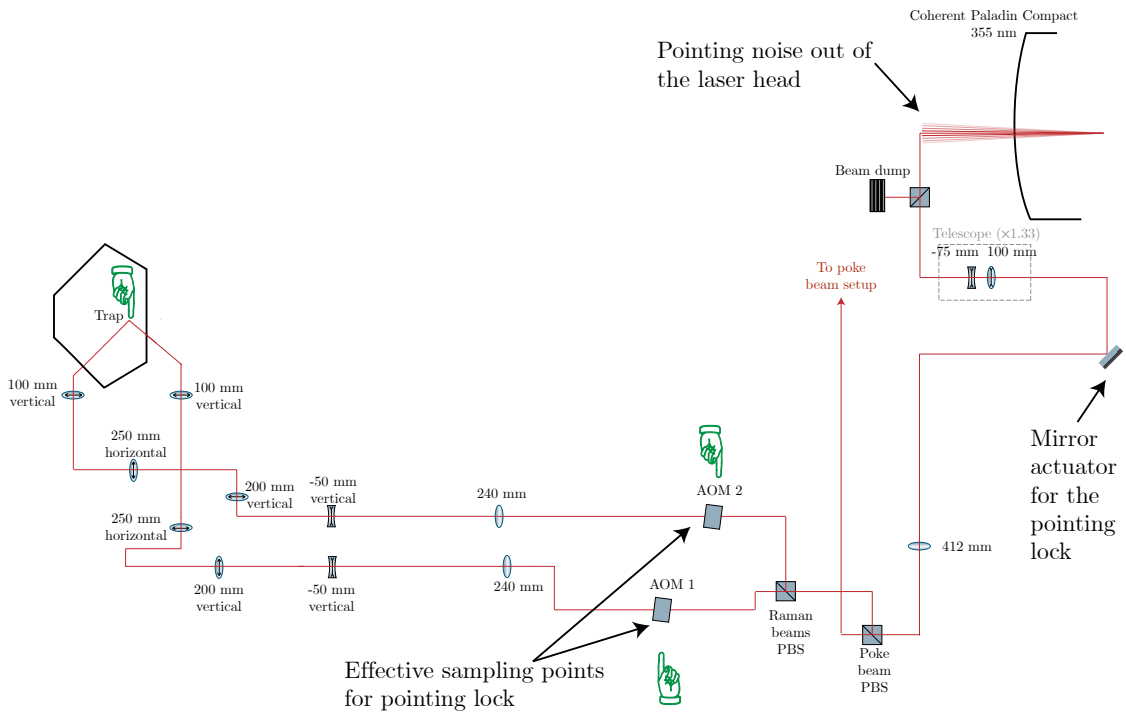

Figure 1.5: **Raman bean path and pointing lock.** Green hands indicate beam foci. The predominantly angular noise out of the laser head is mapped to beam displacement noise at the first focus, which is at the Raman AOMs location. We sample the beam (sampling arm not shown) at this effective location and direct it to a quadrant photodiode (QPD). The QPD is part of a PI stabilization circuit that uses the indicated half-inch diameter mirror as an actuator. The schematic omits elements that are not relevant to the pointing lock, such at the intensity stabilization AOM and mirrors. All lenses are shown.

is ON. In other words, the stabilized beam picks up pointing noise as it propagates past its locking point.

Even though beam tubes are placed sporadically, most of the Raman beam setup is enclosed in a box with apertures for beam entrance and exit. Pointing noise measurements show that the box slightly reduces pointing noise at the AOM location, although not nearly as much as the active locking. I note that the enclosures of the 355 nm path, including the Raman beam path, are far from hermetic. To circumvent the issue that no matter how good a pointing or power lock is, the beams will pick up noise as they propagate downstream, we are considering implementing a sample-and-hold scheme. In this case, Raman light would be turned on for example at the beginning of an experimental sequence at a fixed power, sampled as close to the ion location as possible, and then pointing- and power-corrected to the value that was held during the previous sampling.

1.1.3.2 The BB1 composite pulse scheme

Ion chains with 25 ions are about 60-μm long assuming an axial confinement of $f_{\text{axial}} = 0.34$ MHz. For Raman beams with a horizontal waist of 100 μm, the edge ions are receiving about 90% of the intensity the central ions receive (accounting for the beam–chain angle of 45°), as illustrated in Fig. 1.6. Therefore, a π-time measured with the central ions will not be accurate for the edge ions and will lead to miscalibrated rotations and dephasing across the chain. To counteract this effect, we make use of the BB1 broadband composite pulse scheme [17]. It allows for Bloch

rotations to be done accurately even at the presence of relatively large systematic amplitude errors, i.e. miscalibrated π-times. A perhaps equally important benefit of the BB1 scheme is that it protects us from slow beam intensity noise as well. After all, the method ensuring precise Bloch rotations for slight intensity variations does not care whether these variations are spatial or temporal. Of course, the portion of the noise *during* the rotation itself still affects us.

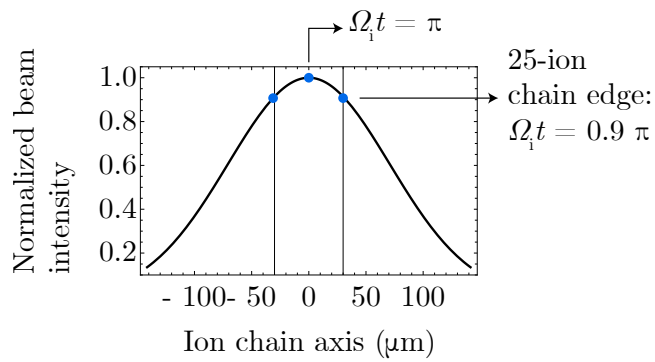

Figure 1.6: **Gaussian beam profile along the chain axis.** For a trap just enough tight to not make a 25-ion chain go to the zig-zag configuration, the edge ions are about 30 μm away from the chain center and the light intensity they see is 0.9 of the maximum at the center. The beam waist is assumed to be 100 μm and the beam at a 45° angle with the chain axis. For axially looser traps that accommodate > 25 ions, the inhomogeneity is amplified, since the 25 central ions will now be even more spread out.

A rotation unitary acting on the Bloch vector ψ and rotating it by an angle θ, around an axis on the XY plane at an angle ϕ with the positive X axis can be written as

$$R_\phi(\theta) = e^{-i\frac{\theta}{2}\sigma_\phi}, \qquad \sigma_\phi = \sigma_x \cos\phi + \sigma_y \sin\phi$$

The actual rotation angle really has a spatial dependence and varies for each ion i: $\theta_{i,\text{act}} = \Omega_i t$, where Ω_i is the Rabi frequency experienced by ion i (see Fig. 1.6). So, it is sensitive to the spatially-inhomogeneous intensity profiles of the Raman lasers.

We instead apply the following BB1 unitary, consisting of 4 sub-rotations:

$$R_4 R_3 R_2 R_{\text{original}} = R_{\phi_{\text{BB1}}}(\pi) R_{3\phi_{\text{BB1}}}(2\pi) R_{\phi_{\text{BB1}}}(\pi) R_{\phi}(\theta) \tag{1.3}$$

where ϕ_{BB1} determines the axis of the three additional rotations, and depends solely on the desired rotation angle θ and not the desired rotation axis:

$$\phi_{\text{BB1}} = \arccos\left(\frac{-\theta}{4\pi}\right)$$

Note that the first rotation to be applied in (1.3) is the original one.

The trick is that if the original rotation angle is not θ, but $\theta + \epsilon$ because of a difference in the Rabi frequency between sites i for example, the error in the combination of all four rotations in (1.3) will be much smaller than ϵ. In theory, a systematic miscalibration of the order of 10% will still yield a a negligible error in the final composite rotation.

We typically use BB1 pulses when the accuracy of the rotation required is high. For a typical scenario where the desired rotation around some axis is π, we would need to measure the π and 2π times with a single ion, and input them in our control software for the durations of the four rotations in (1.3). We find however that an extra step is needed to further improve the BB1 pulses fidelity. This is because, while they can significantly reduce systematic intensity errors, they cannot account for unexpected imperfections in the relative angles of the four rotations. For example, if for some reason the assumed 2π time input for rotation R_3 does not

produce exactly twice the angle of rotations R_2 and R_4, the scheme will not fully work. In other words, all the rotation durations needed might be off by an *overall* factor multiplying all of them, but we cannot tolerate relative discrepancies.

To further elucidate this, assume that the accurate 2π-time for a rotation is $t_{2\pi}$. If we aim to do a π rotation with the BB1 scheme and we use the durations

$$\frac{1}{2}1.05t_{2\pi}, \quad \frac{1}{2}1.05t_{2\pi}, \quad 1.05t_{2\pi}, \quad \frac{1}{2}1.05t_{2\pi}$$

for the rotations R_{original}, R_2, R_3, R_4 respectively, we effectively introduce a global systematic miscalibration of 5%. The BB1 scheme should be able to handle that fine.

- If we input

  $$\frac{1}{2}1.04t_{2\pi}, \quad \frac{1}{2}1.05t_{2\pi}, \quad 1.05t_{2\pi}, \quad \frac{1}{2}1.05t_{2\pi}$$

  instead, then the BB1 scheme will also work, but it will produce accurate rotations by an angle $1.04/1.05\pi$. These should be robust to slow intensity noise for example, or the spatial inhomogeneity of the beam along the chain. However, ...

- if we input

  $$\frac{1}{2}1.05t_{2\pi}, \quad \frac{1}{2}1.05t_{2\pi}, \quad 1.04t_{2\pi}, \quad \frac{1}{2}1.05t_{2\pi},$$

  i.e. a relative miscalibration between the $\pi-$ and 2π-times, then the scheme will not work. What we see in this case, when we do a sequence of multiple BB1 composite π-pulses, is a beatnote instead of a pure oscillation. For slight

miscalibrations, we only see the first part of the beatnote, a slow decay (upper plot of Fig. 1.7)

In practice, when we measure the π- and 2π-times with one ion, it is possible that one of these will be off in its accuracy relative to the other. To address that, we execute sequences of many BB1 π-pulses, like 100, as shown in Fig. 1.7, and tweak the duration of the 2π-pulse by small amounts of the order of 0.5%, until we see no decay in the oscillation envelope.

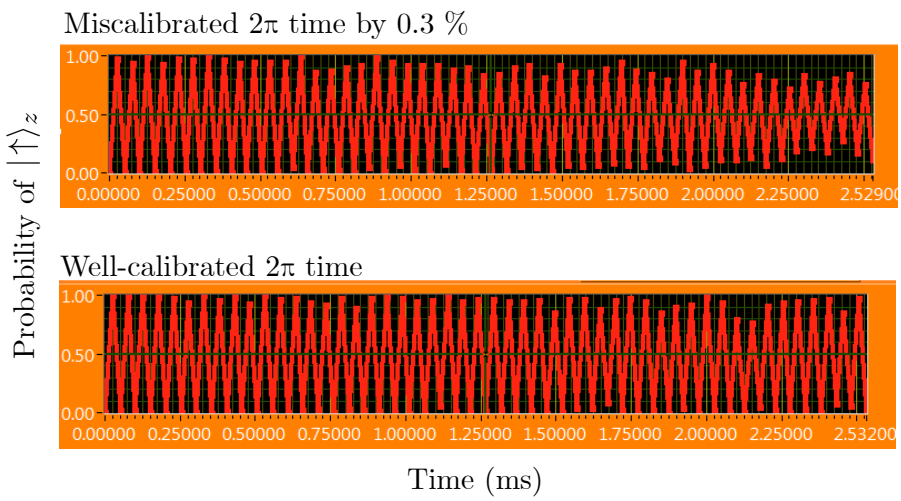

Figure 1.7: **Sub-optimally calibrated and optimally calibrated BB1 pulses.** A miscalibration of the 2π time by a mere 0.3% adds up when we execute 100 π-pulses (plot above).

1.2 The ion trap

To manipulate the ions with laser light, we must spatially confine them so we can address them. Additionally, this has to be done in a manner that minimizes their interaction with their environment, aside from the intended laser-ion interactions. To do that, our lab (QSim) is operating a three-layer linear RF trap [18], which

falls under the category of Paul traps [19]. A combination of static and oscillating electric fields is used to create a potential that confines the ions along all three spatial axes. The confinement along Z is weaker, resulting to a linear arrangement of the ions into a chain along this axis. The spacing between them is not uniform and is determined by the interplay of the repulsive Coulomb force and the trapping potential. Along any axis, the potential near the local minimum, where the ion chain lives, is approximated by a quadratic term $1/2m\omega_q^2 q^2$, where $q = x, y, z$. The secular frequencies ω_q are a measure of how "tight" the potential is along q and typical values are $\omega_x = \omega_y \approx 2\pi \cdot 4.7$ MHz and $\omega_z \in 2\pi \cdot [0.3, 0.6]$ MHz. A schematic of our trap is shown in Fig. 1.8a. The principles of operation of Paul traps have been documented in previous theses and of course, literature, so I will be selectively focusing on specific aspects here. First, I will present an intuitive explanation for how the confinement along two axes, X and Y, is done using a time-dependent field, while still admitting a static confining potential approximation. Second, I will explain and most importantly, visualize our method for ensuring that the two Raman beams impart momentum only along the X axis, while leaving Y undisturbed.

1.2.1 Emergence of the static pseudopotential from the RF drive

The confinement along X and Y is done with an oscillating RF field at a frequency Ω_{RF}, whose effect on the ions is averaged over an oscillation period to yield a static pseudopotential ψ_p. This practice has its own challenges [21], but is a necessity for ions that need to be localized in all three axes: charged particles

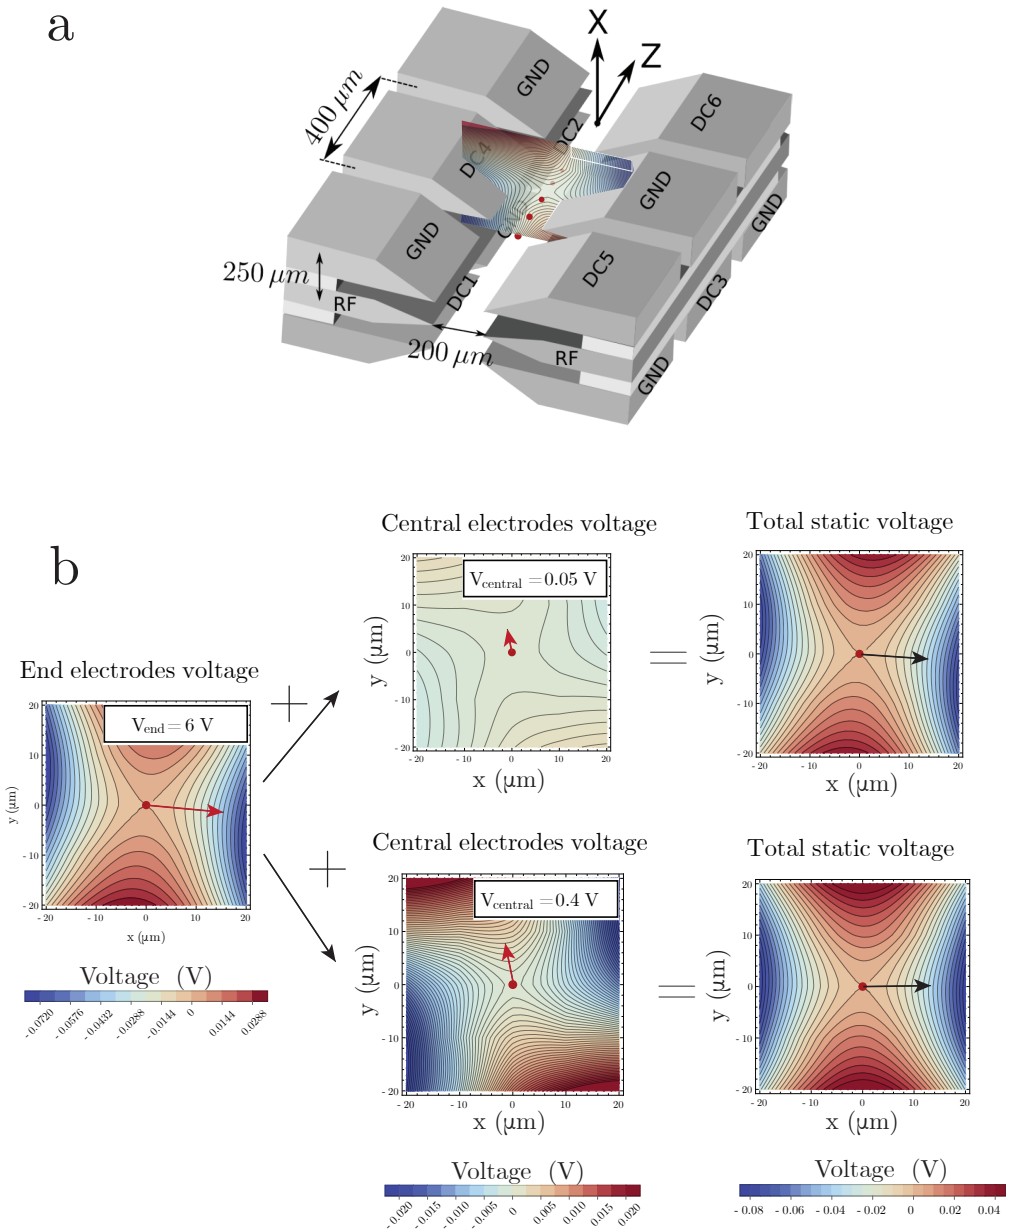

Figure 1.8: **Ion trap schematic and method for rotating the transverse axes X and Y. a.** The 3-layer ion trap in QSim. The Raman beams are along the XZ plane and the ion fluorescence is collected along the X axis. To rotate the XY plane, we must effectively rotate the inset (equipotential lines) around Z. Adapted from [20]. **b.** Two different scenarios for the central electrodes (DC3 and DC4) voltage show how they are used to rotate the trap's XY (transverse) plane. To visualize this rotation, I plot the potential gradient (arrow). The total gradient (black arrows, right column) is the sum of the one created by the end and the central electrode voltages. Each of these gradients points to a fixed direction, therefore by changing their relative magnitudes, we can steer their sum by small angles. For a fixed end electrode voltage of $V_{\text{central}} = 6$ V and two different scenarios of For $V_{\text{central}} = 0.05$ V and 0.4 V, I plot the total gradient, demonstrating its angle control.

cannot be localized by static electric fields alone. Interestingly, this was realized by Earnshaw in 1839 [22], before the scientific community rejects the idea of *ether*, as attested by the assumption of "luminiferous ether" throughout that paper!

This oscillating field is produced by the RF electrodes, shown in Fig. 1.8a. The maximum voltage at the RF electrodes is multiple hundreds of Volts and it results in a "breathing" potential around the ions location; see Fig. 1.9. An ion that sits exactly at the center of this breathing potential, called the *RF null*, will not feel any force. An ion that is slightly displaced from the Rf null though will move under the influence of this force with the $\cos \Omega_{\mathrm{RF}} t$ time dependence, executing periodic motion. This motion has a very small amplitude and is called *micromotion*; it is the less interesting part of ion motion.

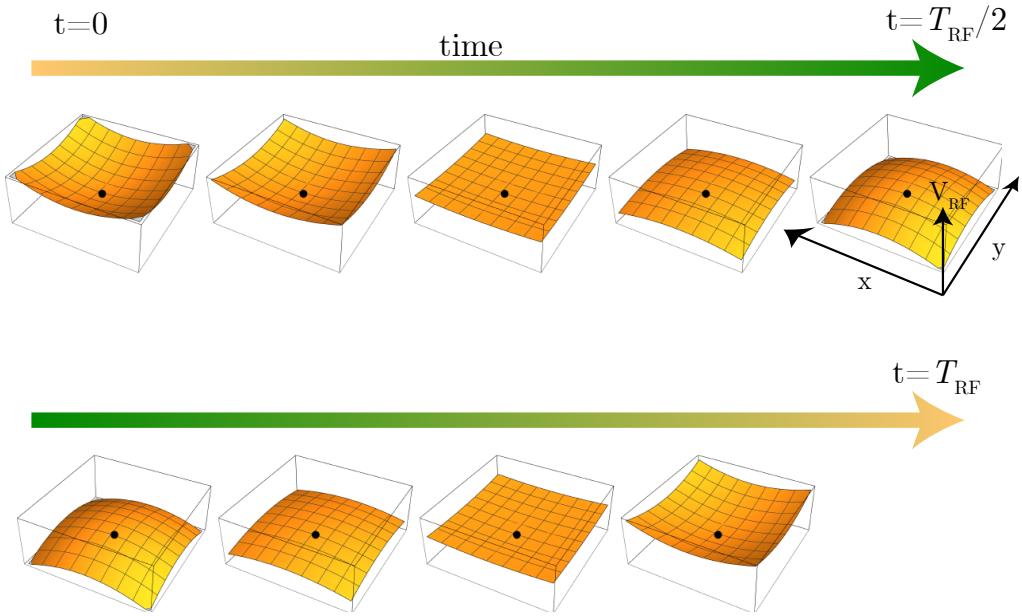

Figure 1.9: **Time-dependent RF potential during one RF period T_{RF}.** Shown is the oscillation of the time-dependent part of the potential along X and Y as experienced by an ion at $Z = 0$, i.e. the middle of the chain. The potential at $t = 0$ was calculated at the trap construction phase using CPO software. The range of X and Y shown is 100 μm. The dot is the ion equilibrium position, i.e. the RF null.

To understand the more interesting motion that arises, called the *secular* motion, let's try and understand what happens if the ion is displaced from the RF null. For simplicity, I will consider the 1D problem, as depicted in Fig. 1.10. In this case, an ion that is very close to the RF null, i.e. the $x = 0$ position, will be subject to the RF force during its oscillation around its average position $\bar{x}$. Since $\bar{x} \approx 0$, the ion's micromotion is symmetric around $x = 0$ and so is the force it experiences: its time average for an RF cycle will be zero. This will not be true if the ion's micromotion is originally centered farther away from the RF null, eg. at a position with $\bar{x} < 0$. In this case, when the ion is at its leftmost position, it experiences a force F_{RF} pointing to the right. When at its rightmost position, F_{RF} points to the left, but with a smaller magnitude. This is exactly because the RF electric field and as a result the experienced force have a *spatial gradient*. Because of this asymmetry, the average RF force during one cycle in this case will not be zero, but point towards the right. Equivalently, if the ion's micromotion is originally centered at the left of the RF null, $\bar{t} > 0$, the time-averaged force will be pointing towards the left.

So far we proved that the time-averaged force $\bar{F}_{\mathrm{RF}}$ during an RF cycle points towards the RF null. Another step is needed to dub $\bar{F}_{\mathrm{RF}}$ a restoring force: it needs to be sufficiently linear to the ion's displacement from the origin. To argue for that, we can look at the RF electric field shown in the top left Fig. 1.10: the farther from $x = 0$ an ion is, the larger in magnitude the force is. Therefore, its acceleration will also increase; and with the acceleration, its micromotion amplitude. Since now the ion is sampling a larger length, the difference of the RF force at its extrema will increase, meaning that the time-averaged force during the RF cycle also increases.

For an electric field that is linear with x, this effect will also be linear to the ion's displacement from the RF null, and therefore we can call $\bar{F}_{\mathrm{RF}}$ a restoring force, giving rise to an effective static potential ψ_p, called the *pseudopotential*.

After these intuitive arguments, one can proceed to calculations and derive the pseudopotential from the time-dependent electric field $E_{\mathrm{RF}} = E_0(x)\cos\Omega_{\mathrm{RF}}t$ [23]:

$$\psi_p(x) = \frac{eE_0^2(\bar{x})}{4m\Omega_{\mathrm{RF}}} \tag{1.4}$$

where ψ_p has units of Volts, e is the elementary charge, and m is the ion mass. As long as $E_0(x) \sim x$, ψ_p will be quadratic. Sufficiently far from the RF null, the electric field $E_0(x)$ will not be linear to the position x any more, meaning that $\psi_p(x)$ is not harmonic any more either.

1.2.2 Minimizing coupling of the Raman beams to the Y motional modes

Our Raman beams are parallel to the optical table and their wavevector difference imparts momentum $\Delta\vec{k}$ to the ions. This momentum "kick" is purely along the X axis if the confining potential's principal axes[5] are optimally aligned to the Raman beams. If they are not, $\Delta\vec{k}$ has a finite projection along the Yaxis, exciting the ion chain's motional modes along Y. This usually happens after the Raman beams have needed to be considerably steered to maximize their intensity at the

[5]The principal axes on a plane are merely the directions of maximum and minimum potential gradient.

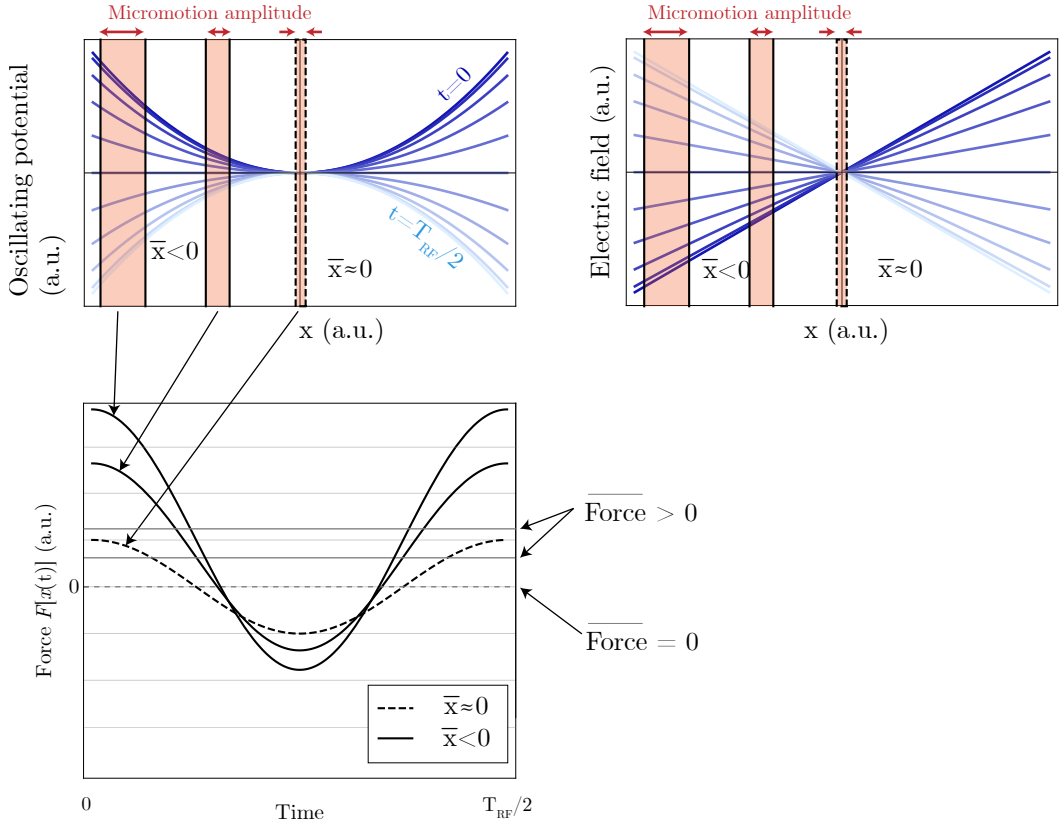

Figure 1.10: **Emergence of the secular motion from the time-dependent electric field gradient. Top row:** the time-dependent potential and electric field E_{RF} as a function of position x. Shades of blue are used to denote different moments during the RF cycle, from $t = 0$ (dark blue) to half-way through, $t = T_{\mathrm{RF}}/2$ (light blue). An ion whose micromotion is centered around the RF null $x = 0$ (dashed rectangle) experiences a force from the RF field that's equal in magnitude at the extrema of its motion; therefore the average force during a single RF cycle is zero. Conversely, an ion whose micromotion is centered left of the RF null (solid rectangles) will be experiencing a larger force towards the right when it is at its leftmost position, than a force towards the left when at its rightmost position. Therefore, the average force during a single RF cycle will point toward the right. This effect is amplified the farther away from the RF null the ion is. This average force gives rise to an effective harmonic oscillator potential and the secular motion of the ion. **Bottom:** The force $F_{\mathrm{RF}}[x(t)] = eE_{\mathrm{RF}}[x(t)]$ from the RF field throughout one RF cycle. Note that this force is the one experienced from the ion during its micromotion, so it is proportional to $eE_{\mathrm{RF}}[x(t)]$ rather than $E_{\mathrm{RF}}[x, t]$ for a fixed position x. For the ion that is roughly at the RF null (dashed), the average force during the RF cycle is $\overline{\mathrm{Force}} = 0$, as stated above. For the ions at the left of the RF null ($\bar{x} < 0$), the average force during the RF cycle is $\overline{\mathrm{Force}} > 0$.

ion location, or when the trap voltages have been changed for some reason. To recover the trap axes-Raman beam alignment, we slightly rotate the trap's XY plane around Z. Since physically rotating the trap chamber is not an option, we change the electrode voltages in a manner that effectively performs this rotation. We call this "minimizing the Y mode", i.e. minimizing the coupling to the Y mode by the Raman beams.

This procedure is illustrated in Fig. 1.8b. While the confinement in the XY plane is provided by the RF electrodes and the pseudopotential they create, we can use the static electrodes to superimpose a weak static potential that adds to ψ_p. By steering the principal axes of this weak potential, we align the Y axis to the $\Delta\vec{k}$ of the Raman beams. The control knob is the voltage of the central electrodes. We typically change it in steps of 0.05 V to 0.1 V, performing a Ramsey scan of the sideband frequency after each step. The duration of the ON time for the Raman beams is chosen so that the main mode (X) shows a full-contrast peak, i.e. it is a π-time for the X mode. We then try to change the central electrode voltage to minimize the other (Y mode) peak.

1.2.3 The vacuum

In order to minimize the ions' interaction with the environment and keep them localized, it is necessary to preserve their surroundings as...empty and boring as possible. This is accomplished by placing the ion trap, i.e. the electrodes structure, inside an ultra-high vacuum (UHV) chamber. Windows in three of its sides provide

optical access so the laser beams can reach the ions and exit the trap, and so that the ions' fluorescence can exit the chamber and be collected by our imaging optics. During the five years of my stay in the lab, the vacuum did not present any issues that required attention of longer than a few hours, and certainly did not require opening it.

The measure used to assess the vacuum is the air pressure in Torr; ours is about 10^{-11} Torr $\approx 10^{-9}$ Pa $\approx 10^{-14}$ atm. From the basics of Thermodynamics, recall that the pressure depends on both the number of gas particles and their kinetic energy. The number of background gas molecules is primarily determined by the pumping system in use: the ion pump and the TiSub pump that was used in the past. The molecule that is less efficiently pumped out is hydrogen H_2, so a common assumption is that most of the pressure is due to hydrogen molecules. As for their kinetic energy, this is determined by the Maxwell-Boltzmann distribution, centered around an average kinetic energy $\langle E_{\text{kin}} = 3/2 k_B T \rangle$ where $T \approx 300$ K and k_B is the Boltzmann constant. During my stay in the lab and for a few years before that, the TiSub pump was not "fired", i.e. heated. Despite this being a recommended periodic procedure, the fear that heating it after a multi-year recess could cause unpredictable vacuum issues kept us from doing so. Besides, ion chain lifetimes were exceptional for short periods of a few weeks at a time, indicating that bad vacuum is likely not the limiting factor for the rest of the time, when the lifetimes are suboptimal.

1.2.3.1 Is your ion chain blinking like Christmas lights?

Two to three times in the past few years we observed that the ion chain was blinking: ions were randomly switching between dark and bright at a fast rate. There are two reasons why this might be happening, and one of those is bad vacuum: if there are many background molecules colliding with the chain, it will be rearranging between different spatial configurations in response to these elastic collisions, appearing to be blinking. The other reason is related to malfunctioning pumping from the $^2D_{3/2}$ state owing to an issue with the 935 nm repump laser and its EOM sidebands.

1.3 Simulating the Ising model and variants

At the heart of our experimental sequences is the part where the Ising-like Hamiltonian is generated. This is done using the Molmer-Sorensen scheme [24], in which we do not apply one Raman beatnote as in the single-ion rotations case, but two at the same time. These "red" and "blue" beatnotes are detuned by δ from the transverse motional sidebands $\pm\omega_{\mathrm{COM}}$ of the carrier frequency ω_{HF}:

$$\begin{aligned} \text{red beatnote: } \omega_{\mathrm{HF}} - \mu &= \omega_{\mathrm{HF}} - \omega_{\mathrm{COM}} - \delta \\ \text{blue beatnote: } \omega_{\mathrm{HF}} + \mu &= \omega_{\mathrm{HF}} + \omega_{\mathrm{COM}} + \delta \end{aligned}$$

In a basic scenario, we apply these beatnotes to a pair of ions both initialized in $|\downarrow\rangle_z$. They will flop from $|\downarrow\downarrow\rangle_z$ to $|\uparrow\uparrow\rangle_z$ and back, and be entangled in-

between, in a state that looks like $\cos J_{12}t \cdot |\downarrow\rangle_z|\downarrow\rangle_z + \sin J_{12}t \cdot |\uparrow\rangle_z|\uparrow\rangle_z$, $J_{12} \in \mathbb{R}$. The underlying basis for this flopping can be intuitively understood with a reasoning somewhat reminiscent of Raman Rabi flopping. The two states $|\downarrow\rangle_z|\downarrow\rangle_z$ and $|\uparrow\rangle_z|\uparrow\rangle_z$ are coupled to each other with the mediation of the states $|\downarrow\rangle_z|\uparrow\rangle_z$ and $|\uparrow\rangle_z|\downarrow\rangle_z$. We choose δ to be large enough so that the beatnote is sufficiently off-resonant from the (red- and blue-sideband) transitions for any motional mode. The role of the two Raman beams in Raman transitions is played here by the two Raman beatnotes, red and blue. The coupling strength J_{12} is

$$J_{12} = \sum_{k=1}^{2} \frac{\eta_{1k}\eta_{2k}\Omega_i\Omega_j}{\mu^2 - \omega_k^2}\omega_k \tag{1.5}$$

where ω_k is the frequency of motional mode k, Ω_i is the carrier Rabi frequency of ion i (at laser intensities that are the geometric mean of the red and blue beatnote intensities), and η_{ik} is the Lamb-Dicke parameter η_k for mode k weighted by the the i-th ion's participation b_{ik} in that mode[6]:

$$\eta_{ik} = b_{ik}\eta_k \tag{1.6}$$

When we have $N > 2$ ions, J_{12} assumes the role of the strength of the interaction between spins i and j in an Ising-like Hamiltonian

$$\sum_{i<j}^{N} J_{ij}\sigma_i^x\sigma_j^x \tag{1.7}$$

[6]Essentially, b_{ik} is the normalized amplitude of classical motion of ion i for mode k, $\sum_k b_{ik}^2 = \sum_i b_{ik}^2 = 1$.

with

$$J_{ij} = \sum_{k=1}^{N} \frac{\eta_{ik}\eta_{jk}\Omega_i\Omega_j}{\mu^2 - \omega_k^2}\omega_k. \tag{1.8}$$

Note that as mentioned in 1.2.2, our Raman beams only couple to the X transverse modes, so the sum above goes up to $k = N$. If we coupled to the Y modes as well, we would need to sum up to $k = 2N$. The coupling element J_{ij} is often approximated by a power law

$$J_{ij} \approx \frac{J_0}{|i-j|^\alpha}, \qquad 0 < \alpha < 3 \tag{1.9}$$

which is a good approximation for $N < 20$ ions, empirically.

Our simulations commonly involve effective magnetic fields in addition to the spin-spin interaction term in (1.7). The generic Hamiltonian we apply can be written in the form of a long-range Ising model with fields:

$$H_{\text{sim}} = \sum_{j<i}^{N} J_{ij}\sigma_i^q\sigma_j^q + B_{q'}\sum_{i}^{N}\sigma_i^{q'} + \sum_{i}^{N} B_i^z\sigma_i^z \tag{1.10}$$

where q, q' are axes on the $X - Y$ plane of the Bloch sphere. In theorists language, the first term is "2-local", which means that the interactions are 2-body (between pairs of spins), even though they act on all pairs of them. I am not going through the full derivation of how (1.10) emerges from the physical laser-ion interaction Hamiltonian, but I will focus instead on how the profile of the J interaction matrix depends on the structure of the chain's motional modes and the detuning of the laser beatnote from them. For the mathematical derivation, reference [25] and the PhD theses [26] and [27] are good references; [28] also nicely illustrates the small

time-dependece of J in small times of the order $1/\delta$.

Generation of the site-dependent field $B_i^z \sigma_i^z$ is afforded by means of a tightly-focused laser beam that can address ion i with minimal crosstalk between neighboring ions [29]; its waist is of the order of 0.5 μm. In addition to the initialization and measurement abilities mentioned earlier, the parameters q, q', J_0, α, $B_{q'}$, and B_i^z are tunable and together with single-spin rotations form a versatile toolkit.

The detuning δ_k of this beatnote from each mode k controls the coupling to that mode, determining its contribution to the resulting Hamiltonian. Since the mode eigenvector b_{ik} has a different profile with regard to i, the detunings δ_k allow for shaping of the interaction matrix J.

1.3.1 The interaction matrix and the motional modes

The profile of the interaction matrix J depends on the mode eigenvector b_{ik} for each mode k and on how strong we couple to that mode, determined by the detuning of the beatnotes from it. To get an idea about these relationships, let's simplify the expression (1.8) for J_{ij}. I will use $\mu^2 = (\omega_k + \delta_k)^2 \approx \omega_k^2 + 2\omega_k\delta_k$ since δ_k/ω_k is of the order of 0.01. I will also neglect the spatial dependence of the beam intensity: $\Omega_i = \Omega \quad \forall i$. We can then write

$$J_{ij} \approx \frac{\Omega^2}{2} \sum_k \frac{\eta_{ik}\eta_{jk}}{\delta_k} \tag{1.11}$$

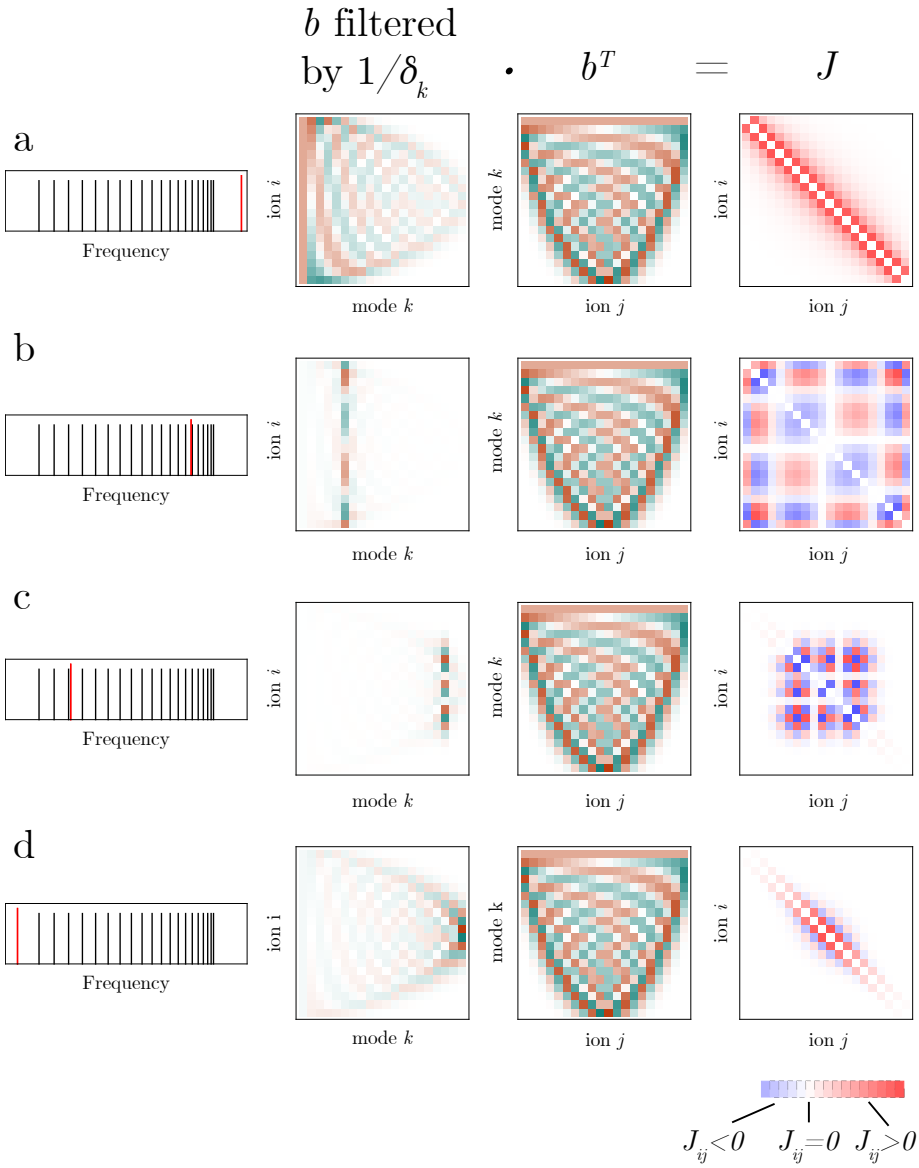

Figure 1.11: **Varying detuning scenarios.** Four different scenarios, $a - d$, for the detuning δ of the Raman beatnote (red line) from the motional modes, shown at the left in each row. The J_{ij} matrix is the product of the two matrices shown in the second and third columns. Our typical case is a. For detunings outside the forest of transverse modes (scenarios a and d), the J matrix looks similar in magnitude, but in d it contains negative (ferromagnetic) couplings in addition to positive ones. In scenario b, the beatnote is tuned close to a mode with nodes on three ions. In the resulting J_{ij} matrix, these spins interact minimally with others. In c, the beatnote is close to a mode where the edge ions participate very little in the motion. The resulting interaction matrix resembles an interacting subset of spins embedded in a system with whom it interacts only weakly, through its boundaries. Color scheme scaling varies for each row.

and using 1.6,

$$J_{ij} \approx \frac{\Omega^2}{2} \sum_k \frac{\eta_k^2 b_{ik} b_{jk}}{\delta_k} \tag{1.12}$$

Given that the bandwidth of the transverse modes in our case is much smaller than their mean, I will neglect the $\eta_k \sim 1/\sqrt{\omega_k}$ dependence and write $\eta_k \equiv \eta = \Delta k/(2m\omega_{\mathrm{COM}})\,\forall k$, where m is the ion mass and Δk the momentum kick from the Raman beams:

$$J_{ij} \approx \frac{(\eta\Omega)^2}{2} \sum_k \frac{b_{ik} b_{jk}}{\delta_k} \tag{1.13}$$

A first observation from the relation above is that the light intensity only enters through the Rabi frequency Ω. This means that the intensity only affects the spin-spin couplings by an overall factor and it does not alter the matrix profile. Since the range α of the interactions is determined by the profile, the light intensity does not matter for the interaction range. More importantly, the role of the mode eigenvectors b_{ik} is clearly shown. I rewrite each term of the sum as

$$\left(\frac{b_{ik}}{\delta_k}\right) \cdot b_{kj}^T \tag{1.14}$$

which points to the following statement: the J_{ij} matrix is a product of two matrices. The first one is the mode eigenvectors b_{ik} "filtered", aka with each column k multiplied by the weight $1/\delta_k$. The second matrix is just the transpose of b_{ik}. This is illustrated in Fig. 1.11, for four different scenarios of detuning δ. Where we place the beatnote determines which motional modes will be coupled to the most. Typically we detune the beatnote higher of the COM mode, which is the highest,

and none of them dominates the sum in (1.14). As another example, if we place the beatnote close to a mode where a few ions are not participating (are stationary), they will not be coupled to other spins in the resulting interaction matrix (second row of Fig. 1.11). Section 3.3 further discusses the spatial inhomogeneity of J_{ij} along the chain axis in our typical detuning scenario.

The second and third interaction profiles of Fig. 1.11 require tuning the Raman beatnote closer to the motional mode frequencies than the first case, which is the typical one, or the last. As discussed in the next section, that will couple to these modes more strongly and will excite phonons with a higher probability if the laser intensity is kept the same. To maintain the same phonon error rate, we would need to decrease the laser intensity I by the same relative factor as the detuning, suffering a decrease in the overall interaction strength that ends up being the same factor as the detuning reduction. This estimate only considers coupling to the phonon mode that is closest to the beatnote; for an accurate comparison, one would need to sum the contributions for all modes.

As an example, the smallest detuning for the third case of Fig. 1.11 is about 15 times smaller than the one of the first case. If δ, I, J_{ij} are the smallest detuning, intensity, and interaction matrix for the first case, and δ', I', J'_{ij} are the same quantities for the third case, we have

$$\delta' \approx \kappa\delta, \qquad I' = \kappa I, \qquad J'_{ij} = \kappa J_{ij}$$

where $\kappa = 1/15 \approx 0.07$. In words, to keep the phonon errors the same, we would end

up with a 15-fold reduced interaction, requiring sequences to be 15 times longer to observe the same physical effect, i.e. keep the product $J_{ij} \cdot \text{time}$ the same. Needless to say, this decision must consider the experimental coherence times, limiting the duration that interactions can be meaningfully applied. To mitigate this effect, we can decide to tolerate a phonon error rate s times more than the original. In this case,

$$I' = \kappa\sqrt{s}I, \qquad J'_{ij} = \kappa\sqrt{s}J_{ij}.$$

1.4 Approximations and error sources

Alas, the ideal evolution of the spins under (1.10) is tainted by a medley of caveats. Approximations in deriving (1.10) from the physical Hamiltonian as well as laboratory noise warrant our attention, if anything to rank their impact on our simulations and prioritize some of them for addressing. To gain an overview of the relevant causes of infidelity, the reader is encouraged to look at references [2, 21, 30] or PhD theses [26, 28, 31]. What follows is a limited listing of some of these caveats. All of them have already been described theoretically in the literature, but quantifying their impact in each of our simulations is a work in progress.

Fluctuating AC Stark shifts seem to be a primary suspect of deviations from ideal evolution according to our current understanding. Essentially, the Raman beams off-resonantly couple to atomic transitions and shift the splitting between the qubit states $|\downarrow\rangle_z$ and $|\uparrow\rangle_z$. This results in terms that correspond to fields $\sim \sigma_z$, adding to the ideal evolution. Since these AC Stark shifts depend on a variety of

factors that are noisy, they are noisy themselves, and so are these terms. Chapter 2 discusses this noise further.

Secondly, as mentioned, the Raman beams off-resonantly couple to the motional modes of the ion chain. Residual excitation of phonons is another source of errors in our simulation. To understand those, we look at the time evolution operator $U(t)$ for the physical Hamiltonian [2, 25]:

$$U(t) = e^{\sum_i \hat{\zeta}_i(t)\sigma_x^i - i\sum_{i<j}\chi_{ij}(t)\sigma_x^i\sigma_x^j}. \tag{1.15}$$

The term involving

$$\hat{\zeta}_i(t) = \sum_k^N \left[\alpha_{ik}(t)a_k^\dagger - \alpha_{ik}^*(t)a_k\right]$$

includes Pauli operators σ_x^i acting on the spin part of the wavefunction and ladder operators a_k, $a_k^\dagger$ acting on its motional part[7]. It is similar to a displacement operator $D(\alpha_k)$, acting on the initial motional state of zero quanta for mode k and transforming it to a coherent state with an average number of $\bar{n}_k = |\alpha_k|^2$ quanta. This term entangles the spin and the motional degrees of freedom and can have complicated effects on our simulations [32]. In QSim and for transverse fields in (1.10) that are much smaller that the nearest-neighbor interactions J_0, most of these effects can be ignored and this term just causes erroneous spin-flips with a small probability p_{flip} to first-order. See 3.1 for more details and a quantitative estimate of p_{flip}.

Another noise source is fluctuations of the magnitude of the spin-spin inter-

[7]The careful reader will notice that $U(t)$ features $\chi_{ij}(t)$, but not J_{ij}. Because of what is called the slow-regime approximation, for times much longer than the inverse motional detunings $1/\delta_k$, $\chi_{ij}(t) \sim J_{ij} \cdot t$.

actions J_0 and fields, if present, because of beam intensity fluctuations at the ion location. The spectrum of this noise is concentrated in frequencies much lower than we were susceptible for the prethermal time crystal experiment presented in Chapter 4 and for this reason I will not focus on it. The interested reader though can easily model it using insights from the Stark shift noise discussion in Chapter 2.

Finally, in the recent years we have needed to address two additional issues. First, as the size of the ion chains used in our simulations increases, the profile of the interaction matrix J_{ij} deviates from a power law. For the prethermal time crystal experiment we remedied this issue by departing from the power law assumption altogether, and working with the raw J_{ij} itself. Whether that Hamiltonian hosted a phase transition or not had to be explored with numerical calculations, instead of deduced from the exponent of a power law. For visualization of J_{ij} matrices, including the extra issue of their inhomogeneity along the chain, see 3.3. The second issue that warranted attention was decoherence caused by the finite duration of our simulations and the resulting spectral "spreading" of these pulses. The shorter the quench duration of H_{sim}, the broader the spectral decomposition of the Raman beatnote, and the larger its effect on exciting the adjacent in frequency motional modes. We found that using a shaping envelope for our MS pulses greatly decreased this effect; see 3.2 for more.

Chapter 2: Stark shifts, their noise, and its causes

The goal of this section is to build intuition and provide insights on the primary source of decoherence for many of our simulations: fluctuating AC Stark shifts. Theoretically calculating expected values from experimental parameters is tedious; I include the derivation in Appendix A. However, when it comes to estimating their fluctuations based on lab noise, scaling arguments are sufficient and all that is needed is an understanding of how the Stark shift depends on quantities like the intensity of the light. I assume basic knowledge of Raman transitions and the Λ scheme used in the lab to couple the qubit states $|\downarrow\rangle_z$ and $|\uparrow\rangle_z$ through the adiabatically eliminated states of the $P_{1/2}$ and $P_{3/2}$ manifolds.

First, I will motivate studying σ_z noise as a relevant factor in our simulations, and show that Stark shifts manifest as σ_z terms. Afterwards, some background

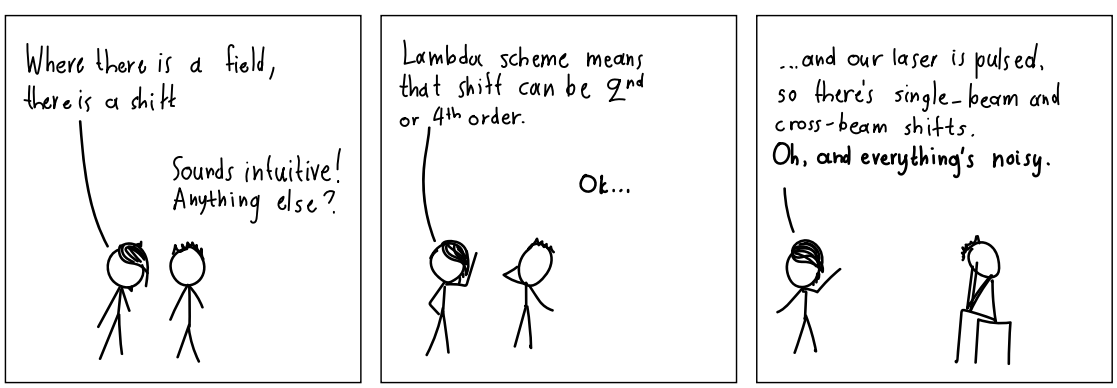

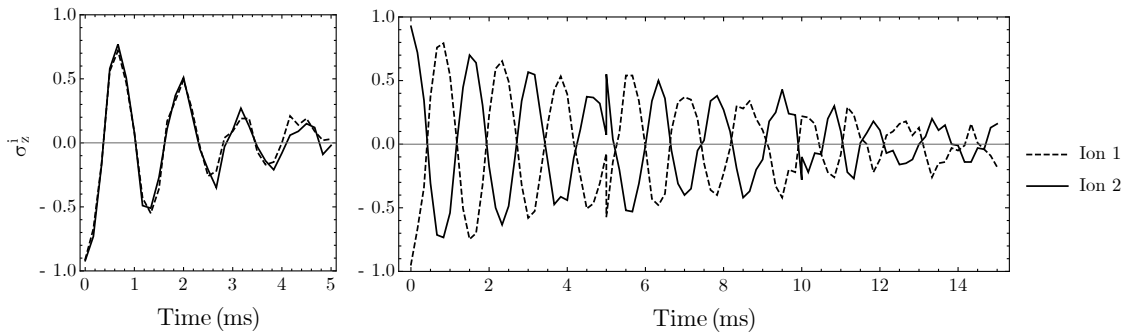

Figure 2.1: The MS coherence time is ~ 3 times longer when the spins initially point to opposite directions in Z.

information on Stark shifts will ready the reader to understand how they depend on physical parameters and how they are categorized. Lastly, I will go over Stark shift noise, presenting the simple model for our lab based on recent measurements.

2.1 Pointing the finger to Stark shift noise

Why do we care about Stark shift noise? On June 2019 we saw compelling evidence that σ_z noise is a major decoherence factor during application of the MS scheme. To show that, the MS scheme was applied to two different initial states with two ions: $|\downarrow\rangle_z|\downarrow\rangle_z$ and $|\downarrow\rangle_z|\uparrow\rangle_z$. In the absence of noise we expect similar results for both these states under the MS evolution [24].

However, as shown in Fig. 2.1, the $|\downarrow\rangle_z|\uparrow\rangle_z$ initial state had a coherence time three times longer than the $|\downarrow\rangle_z|\downarrow\rangle_z$ case. This can be explained if there is a noisy effective magnetic field $B_z(t)\sigma_z$ along Z. To show that, let's look at the time

evolution under this term for infinitesimal time dt for the two initial states:

$$e^{-iB_z(t)\mathrm{d}t\sigma_z}|\downarrow\rangle_z|\uparrow\rangle_z = e^{-iB_z(t)\mathrm{d}t(-1+1)}|\downarrow\rangle_z|\uparrow\rangle_z = |\downarrow\rangle_z|\uparrow\rangle_z \tag{2.1}$$

$$e^{-iB_z(t)\mathrm{d}t\sigma_z}|\downarrow\rangle_z|\downarrow\rangle_z = e^{-iB_z(t)\mathrm{d}t(-1-1)}|\downarrow\rangle_z|\downarrow\rangle_z = e^{-2iB_z(t)\mathrm{d}t}|\downarrow\rangle_z|\downarrow\rangle_z \tag{2.2}$$

We see that for the $|\downarrow\rangle_z|\uparrow\rangle_z$ state, the σ_z noise has no effect, whereas for the $|\downarrow\rangle_z|\uparrow\rangle_z$ state, it adds a phase of $2B_z(t)\mathrm{d}t$. If this noise was the only Hamiltonian term, then this overall phase factor could not demonstrate itself in a measurement of course. But in a realistic scenario with other non-commuting terms present, eg. $\sigma_x^i\sigma_x^j$, this noise contribution is measurable. When different runs are averaged together, each of them has acquired a different phase because $B_z(t)$ is noisy, leading to dephasing.

To strengthen the suspicion that the noise is along Z, we repeat the same experiment for initial states along the Y axis, $|\downarrow\rangle_y|\uparrow\rangle_y$ and $|\downarrow\rangle_y|\downarrow\rangle_y$. Both initial states along Y are expected to decohere comparably fast to $|\downarrow\rangle_z|\downarrow\rangle_z$ and much faster than $|\downarrow\rangle_z|\uparrow\rangle_z$—and they do, as shown in Fig. 2.2:

$$e^{-iB_z(t)\mathrm{d}t\sigma_z}|\downarrow\rangle_y|\uparrow\rangle_y = e^{-iB_z(t)\mathrm{d}t}|\uparrow\rangle_y|\downarrow\rangle_y \tag{2.3}$$

$$e^{-iB_z(t)\mathrm{d}t\sigma_z}|\downarrow\rangle_y|\downarrow\rangle_y = e^{-iB_z(t)\mathrm{d}t}|\uparrow\rangle_y|\uparrow\rangle_y. \tag{2.4}$$

Interstingly, tests performed more recently, in spring 2021, imply that while present, σ_z noise does not fully explain the decay in the MS evolution of Fig. 2.1. In theory, we should be able to echo out such noise, if π-pulses are interleaved with the MS evolution, regardless of the initial state. The interval between the π-pulses

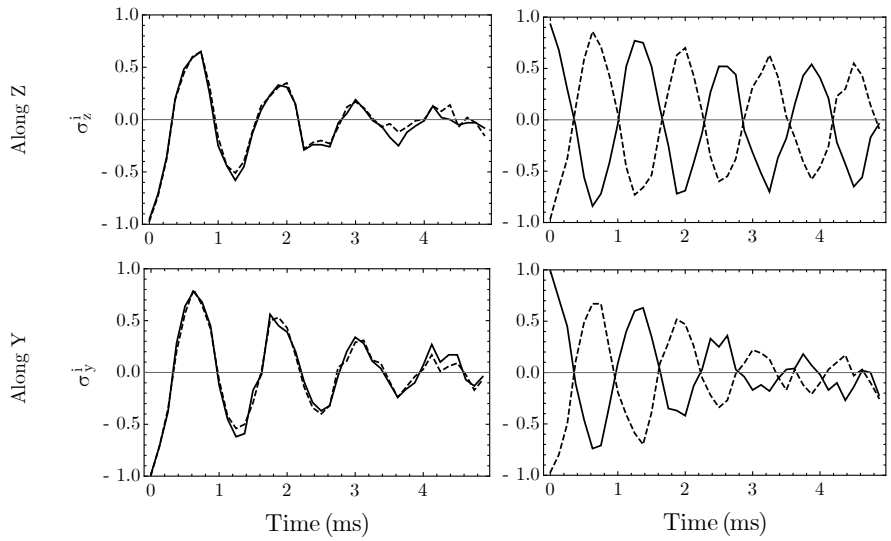

Figure 2.2: Four different initial states strengthen the claim that there is significant σ_z noise. Only the $|\downarrow\rangle_z|\uparrow\rangle_z$ state shows a distinctly longer coherence time.

determines the spectrum of σ_z noise that is echoed out. Indeed, starting from two spins both in $|\downarrow\rangle_z$, the decay time of the MS evolution increased with more frequent π-pulses. However, it never compared to the high values that the $|\uparrow\rangle_z|\downarrow\rangle_z$ initial state had.

2.2 Stark shifts manifest as a σ_z term

After establishing the presence of strong σ_z noise, let's argue that Stark shift noise *could* be its source. Stark shifts pertain to the "physical domain". In the "simulation domain", the Stark shift corresponds to an effective magnetic field along the Z axis. Therefore, we commonly state that a noisy Stark shift corresponds to a noisy σ_z field for the qubits, invoking the Bloch vector picture.

Intuitive explanation. We can think in terms of rotating frames. The qubit's clock is rotating with a frequency of $\omega_{\mathrm{HF}} \approx 12.6$ GHz, and the laser beatnote clock

ideally rotates with the same frequency (for a carrier transition). Any Bloch vector appearing stationary in one of them will also appear stationary in the other.

Now assume that the qubit frequency is Stark-shifted by ΔE. A Bloch vector stationary in the laser beatnote frame will appear to be rotating with frequency ΔE at the qubit frame—as if it's Larmor precessing because of an effective magnetic field! Furthermore, if ΔE fluctuates in time, this effective magnetic field term in the simulation Hamiltonian will also fluctuate.

More detailed explanation. One can use effective Hamiltonian theory to show the connection between a Stark shift and a σ_z term. Let's assume a general form for a Hamiltonian term H_I that describes interactions consisting of one or more harmonic terms:

$$H_I(t) = \sum_{n=1}^{L} \hat{h}_n e^{-i\omega_n t} + h.c. \tag{2.5}$$

where L is the total number of different harmonic terms making up H_I with $0 < \omega_1 \leq \omega_2 \leq ... \leq \omega_L$. The effective Hamiltonian, where all processes oscillating at frequency ω_1 and above are assumed to average out to zero, is [33]:

$$H_{\text{eff}} = \sum_{m,n=1}^{L} \frac{1}{\bar{\omega}_{mn}} \left[\hat{h}_m^\dagger, \hat{h}_n\right] e^{i(\omega_m - \omega_n)t} \tag{2.6}$$

where

$$\frac{1}{\bar{\omega}_{mn}} = \frac{1}{2}\left(\frac{1}{\omega_m} + \frac{1}{\omega_n}\right) \tag{2.7}$$

The important observation is that an interaction Hamiltonian H_I coupling 2 levels $|1\rangle$ and $|2\rangle$ of a 2-level system—as is our case—will include a term propor-

tional to $|1\rangle\langle 2|$ and its hermitian conjugate, $\hat{h} \sim |1\rangle\langle 2|$, and therefore its effective Hamiltonian will be proportional to the commutator

$$H_{\text{eff}} \sim \left[|2\rangle\langle 1|, |1\rangle\langle 2|\right] = |2\rangle\langle 2| - |1\rangle\langle 1| \sim \sigma_z \tag{2.8}$$

2.3 AC Stark shifts and their calculation

A Stark shift is the change in an electronic energy level because of the presence of an external electric field. Classically, the field is pushing the positive nucleus and the negative electron cloud toward opposite directions, overall increasing the potential energy of the atom. In this thesis, all such electric fields are oscillating and not static, and they cause *AC Stark shifts* on an energy level.

It is useful to make a connection between Rabi flopping and Stark shifts. If the light polarization allows for Rabi flopping between states $|i\rangle$ and $|j\rangle$, the same light tuned off-resonance will cause a Stark shift on these states. If the polarization prohibits coupling, there can not be a Stark shift either. The single, simple formula worth remembering relates the Stark shift to the Rabi frequency that the same light would cause *if* on-resonance with the atomic transition [34]:

$$\text{Stark shift} = \frac{(\text{Rabi frequency})^2}{4 \cdot \text{Detuning}} \tag{2.9}$$

Equation (2.9) holds for both 2nd and 4th order Stark shifts. For 2nd order ones, the Rabi frequency is the single-photon Rabi frequency usually denoted by g_0 or g in the literature and the detuning is between the laser frequency ($c/355$ nm in

our case) from the relevant P state, typically $\Delta = 33$ THz or 66 THz. For 4th order Stark shifts, the Rabi frequency is the 2-photon Raman Rabi frequency and the detuning is between the corresponding RF beatnote and the atomic transition between S levels.

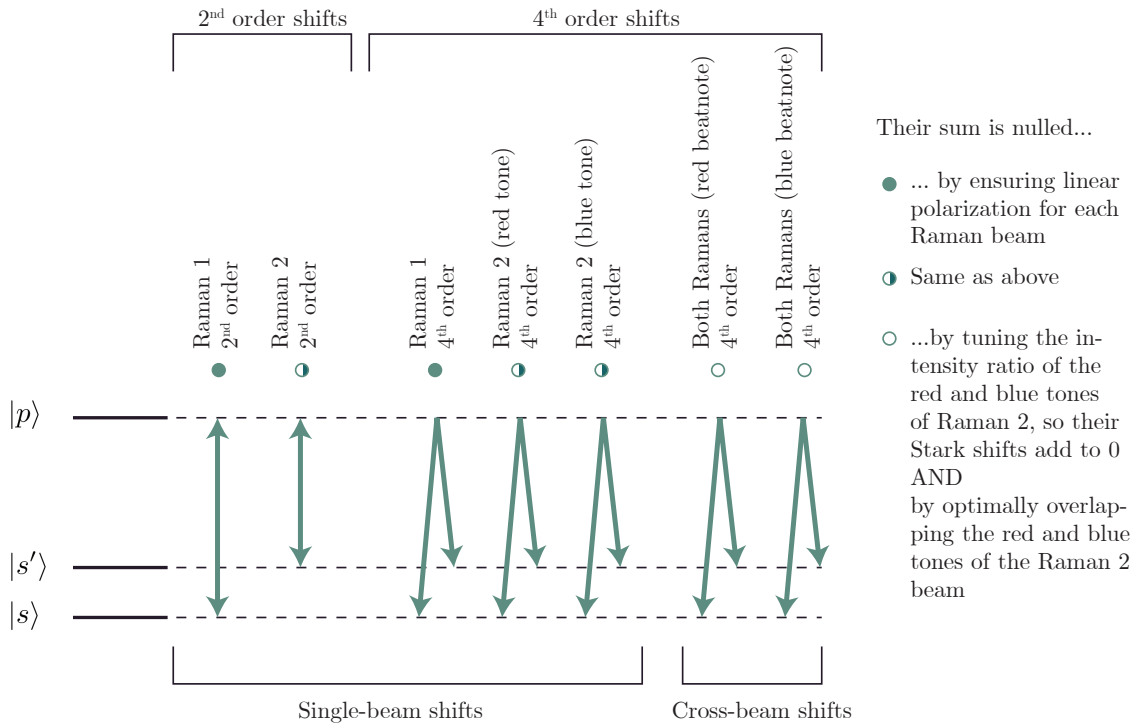

Figure 2.3: **AC Stark shift types**. Schematic of Stark shifts on a state of interest $|s\rangle$ at the S manifold, because of off-resonant light coupling it to a state $|p\rangle$ in the P manifold and/or a state $|s'\rangle$ in the S manifold. Each of Raman 1 and 2 induces a 2nd order Stark shift. During the Ising simulation, Raman 2 has two tones, red and blue, which interfere with Raman 1 to create red and blue beatnotes. For the 2nd Stark shift, we do not need to take into account the red and blue tones of Raman 2, since their frequency difference of $\sim 2f_{\text{COM}}$ is much smaller than the detuning of the light from the transition, $\Delta \approx 33$ or 66 THz. The single-beam 4th order Stark shift is much larger than the single-beam 2nd order. To proceed to calculations, one must further take into account all the states $|p\rangle$ of the P manifold that light can couple with the the $|s\rangle$, $|s'\rangle$ in question.

Note that Stark shifts typically enter our quantum simulation in two ways. First, the single-beam 4th order Stark shift from a tightly focused beam ("poke

beam") is used to engineer programmable local effective magnetic fields [29] . This is a useful, desirable effect, instrumental to the breadth of usage of our quantum simulator. Second, fluctuating cross-beam Stark shifts during the MS scheme manifest themselves as noisy σ_z terms. To our most updated understanding, this σ_z noise is a significant source of decoherence in most of our experiments, and this is the focus of this section.

To calculate the Stark shifts for the qubit states of $^{171}\mathrm{Yb}^+$ from equation (2.9), there are a few steps one needs to do. For a detailed derivation, see Appendix A. These steps consist of accounting for the polarization mix of the Raman beams (potentially a mix of σ^-, π, and σ^+), multiple intermediate P states, multiple S states that can potentially couple to a qubit state via Raman transitions, the fact that our laser beams are not CW, but optical combs, the existence of two optical combs (red and blue) in one Raman beam for our application of the MS scheme, and finally remembering that the 4th order Stark shift measured in the lab is differential, i.e. the difference between the Stark shift of each qubit state.

The jungle of Stark shifts for a qubit state is shown in Fig. 2.3. In theory, they should all sum to zero for QSim's linear Raman beam polarizations, perpendicular to each other (lin$\perp$lin). However, some of them are nulled easier and more reliably than others. Notably, all single-beam Stark shifts are nulled by ensuring the lin$\perp$lin configuration. Measured polarization noise is small (see 2.4.2) with a non-appreciable result in single-beam Stark shift noise. In contrast, nulling the cross-beam shifts from the red and blue beatnotes relies on optimal overlap of the red and blue tones of the Raman 2 beam and fine-tuning of their relative powers;

both these conditions are especially susceptible to noise.

The single-beam and cross-beam Stark shifts for state i take the simple forms shown in Table. 2.1. They are expressed as a product where the light's characteristics are fairly decoupled; these are the intensity (determining Ω_0), the polarization[1] (determining P_{ij}), and the frequency (determining the K_{ij} factors). All these are assumed to be measured at the ion location. The index i in that table denotes the state i whose Stark shift we are calculating, and the index j denote the state j that i couples to via the laser field. Fig. 2.4 shows the Raman beam configuration in QSim and the resulting nulling of the 4th order cross-beam Stark shift. When running experiments, we tune the relative powers of the red and blue tones of Raman 2 to ensure that their Stark shifts add to zero at the center of the ion chain. In addition, every few weeks, especially when optics in the Raman 2 path have been tweaked, we minimize the Stark shift gradient along the ion chain by optimizing the spatial overlap of the red and the blue tones.

Table 2.1: Summary table of Stark shift-related formulas

Single-beam Stark shift: state i, $\Delta E^{(4)}_{\text{single},i}$ Cross-beam Stark shift: state i, $\Delta E^{(4)}_{\text{cross},i}$	$\frac{1}{4}\Omega_0^2 \sum_j P_{ij}^2 K_{\text{single},ij}$ $\frac{1}{4}\Omega_0^2 \sum_j P_{ij}^2 K_{\text{cross},ij}(\delta\omega_{\text{comb}})$
Base Raman Rabi frequency Ω_0	$\frac{1}{6}\left(\frac{g^2_{P_{1/2},0}}{\Delta} + \frac{g^2_{P_{3/2},0}}{\omega_F - \Delta}\right)$
Polarization factor P_{ij} for states i and j	See Appendix A.4
$K_{\text{single},ij}$	$\sum_k \left(\frac{k\omega_{\text{rep}}\tau}{2}\right)^2 \text{csch}^2 \frac{k\omega_{\text{rep}}\tau}{2} \cdot \frac{1}{k\omega_{\text{rep}} - \omega_{ij}}$
$K_{\text{cross},ij}(\delta\omega_{\text{comb}})$	$\sum_k \left(\frac{k\omega_{\text{rep}}\tau}{2}\right)^2 \text{csch}^2 \frac{k\omega_{\text{rep}}\tau}{2} \cdot \left(\frac{1}{k\omega_{\text{rep}} - \omega_{ij} + \delta\omega_{\text{comb}}} + \frac{1}{k\omega_{\text{rep}} - \omega_{ij} - \delta\omega_{\text{comb}}}\right)$

[1] For explanation of symbols see Appendix A

[1] which includes the relationship of the light polarization with the quantizing magnetic field

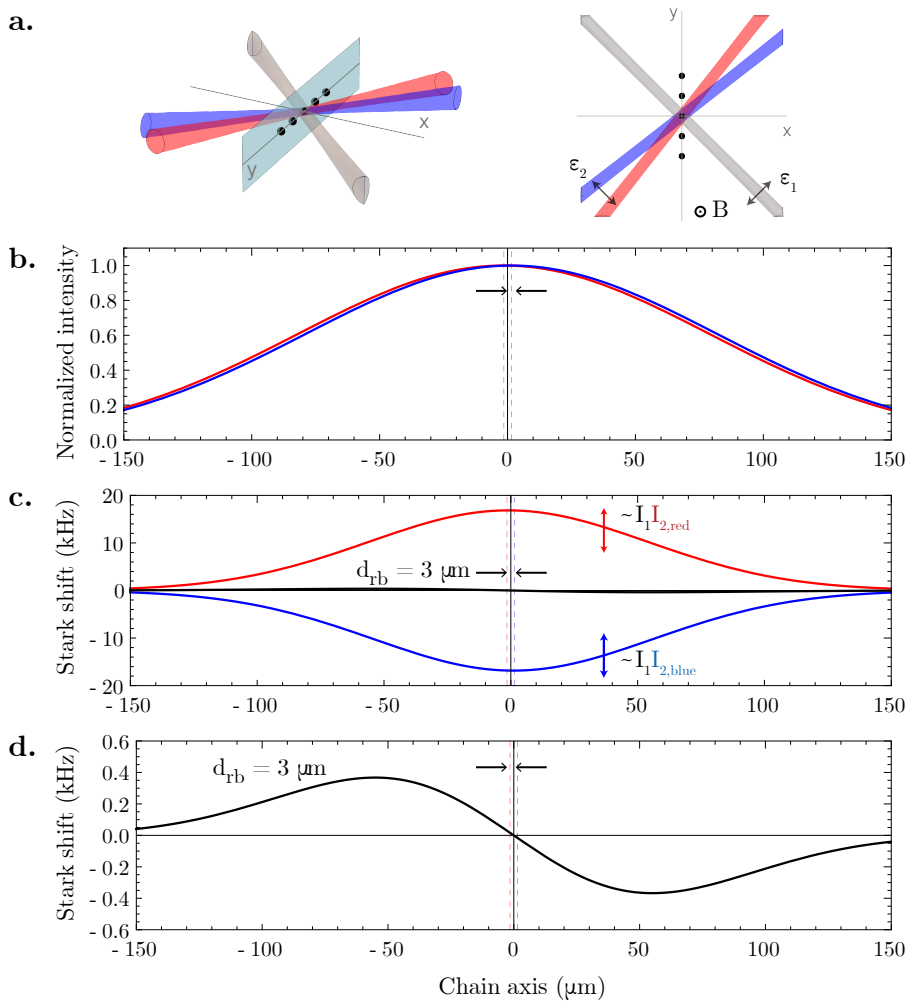

Figure 2.4: **Schematic of the two Raman beams and the Stark shift gradient. a.** Raman 1 (gray) contains only 1 frequency tone. Raman 2 (blue and red) contains 2 tones, which converge at the ion location at a small angle (exaggerated here). The polarizations $\vec{\varepsilon}_i$ and the quantizing magnetic field $\vec{B}$ show the lin⊥lin configuration. **b.** Intensity profile of the red and blue beatnotes at the ion location. The two tones travel slightly separated in space and ideally converge at the ions location. Realistically, our typical tolerance for the distance between the beam centers along the ion chain axis, d_{rb}, is around 3 μm. This is very small compared to the beam waist of about 100 μm. **c.** Each of these tones interferes with the optical comb of Raman 1 to give a red and a blue beatnote. Each of these beatnotes causes a cross-beam 4th order Stark shift with opposite sign. Noise in the Raman 1 or 2 intensity causes each of the red and blue curves to jitter ("breathe") in the vertical direction, as the vertical arrows show. To get the total cross-beam Stark shift, we add the red and blue contributions (black curve). **d.** Zooming in vertically to better discern the black curve (total cross-beam Stark shift) profile, we see that a Stark shift gradient is formed in the region of the ion chain, the center few tens of μm.

2.3.1 Why are the red and blue SB Rabi frequencies different when the Stark shift is nulled?

Part of our calibrations, done a few times per day during data taking, is nulling the sum of the cross-beam Stark shifts of the red and the blue beatnotes. To do that, we tweak the intensities of the red and the blue tones of Raman 2—see Fig. 2.4 for a schematic. We always find that when the Stark shift is nulled, the light power of the blue tone is higher than the red. This is measured by the on-resonance sideband Rabi frequencies, equal to[2] $\eta_{\mathrm{COM}}\Omega_{\mathrm{red}}$ and $\eta_{\mathrm{COM}}\Omega_{\mathrm{blue}}$, where η_{COM} is the Lamb-Dicke parameter for the COM transverse mode and $\Omega_{\mathrm{red}}, \Omega_{\mathrm{blue}}$ the carrier Rabi frequency scaled to the power of each beatnote. The reason that Ω_{blue} has to be larger than Ω_{red} is largely irrelevant to the AOM efficiencies at the different RF drive frequencies for the red and the blue sidebands.

Mathematically speaking, the answer is in the comb factors $K_{\mathrm{cross},ij}$, at Table 2.1. For each beatnote, the frequency difference between the two Raman beams/optical combs is

$$\delta\omega_{\mathrm{comb}}^{\mathrm{red}} = \omega_A - \mu \tag{2.10}$$

$$\delta\omega_{\mathrm{comb}}^{\mathrm{blue}} = \omega_A + \mu \tag{2.11}$$

with $\omega_A = \mod(\omega_{\mathrm{HF}}/\omega_{\mathrm{rep}}) \approx 29.7$ MHz and μ the total detuning of the beatnote from the carrier transition. Plugging (2.10–2.11) in the factors $K_{\mathrm{cross},ij}$ (for $|i\rangle =$

[2] Assuming we start from the 1 or 0 phonon number states for the red and the blue respectively.

$|\downarrow\rangle_z, |j\rangle = |\uparrow\rangle_z$), for typical trap frequencies, we get

$$\left|\frac{K^{\text{red}}}{K^{\text{blue}}}\right| \approx 1.21 \tag{2.12}$$

To make the red-beatnote Stark shift equal in magnitude to the blue-beatnote one, the square of the carrier frequency $(\Omega_{\text{red}})^2$ (see first box of Table 2.1) has to compensate for that factor of ~ 1.21:

$$\frac{\Omega_{\text{red}}}{\Omega_{\text{blue}}} \approx \frac{1}{\sqrt{1.21}} \approx 0.91 \tag{2.13}$$

which means that the sideband Rabi frequencies have to have the same ratio:

$$\frac{\eta_{\text{COM}}\Omega_{\text{red}}}{\eta_{\text{COM}}\Omega_{\text{blue}}} \approx 0.91. \tag{2.14}$$

Fig. 2.5 visualizes this mathematical explanation. The gist is that the red and the blue beatnotes creating the red and the blue Stark shift are different, i.e. the frequencies of their comb teeth are not identical sets of numbers. They might look similar, but the detunings of the teeth of each beatnote from the related transition, ω_{HF}, are a different set for each case. Nominally, the few smallest in magnitude detunings from ω_{HF} are shown in Table 2.2.

Table 2.2: Smallest detunings of the red and blue beatnote teeth from the hyperfine frequency ω_{HF}. Only the smallest detuning matches in magnitude (μ).

Red beatnote	...	$2\omega_A - \mu - \omega_{\text{rep}}$	μ	$2\omega_A - \mu$	$\mu + \omega_{\text{rep}}$	...
Blue beatnote	...	$2\omega_A + \mu - \omega_{\text{rep}}$	$-\mu$	$2\omega_A + \mu$	$-\mu + \omega_{\text{rep}}$	...

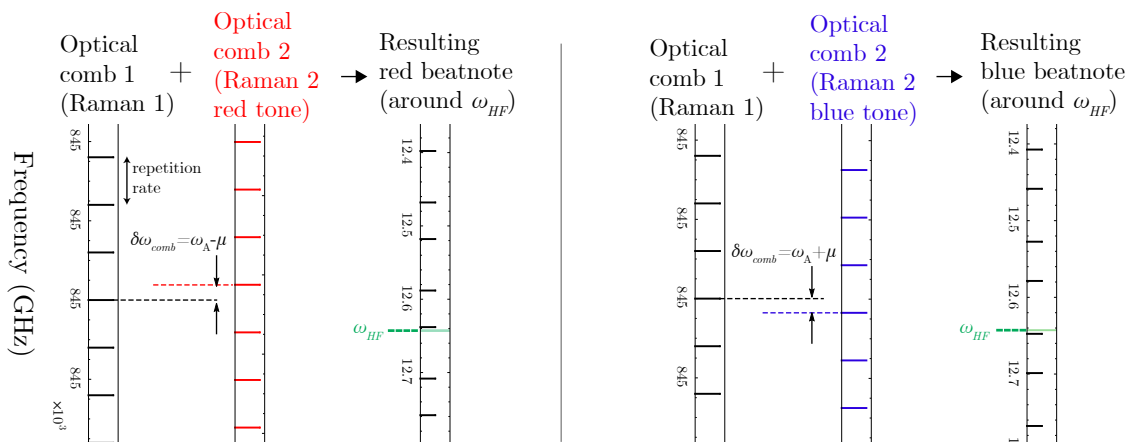

Figure 2.5: **Illustration of the red and the blue beatnotes creating the cross-beam Stark shifts.** Left: Creation of the red beatnote from the interference of the optical combs of Raman 1 and Raman 2–red tone. Right: Creation of the blue beatnote from the interference of the optical combs of Raman 1 and Raman 2–blue tone. The Stark shift for any of the resulting beatnotes is related to how all of its comb teeth are detuned from ω_{HF}, the transition related to this Stark shift (other transitions involving the Zeeman levels of the $S_{1/2}$ manifold are suppressed because of the lin⊥lin polarization scheme—see Appendix A for details). These sets of detunings might look similar for the red and the blue cases, but they are not identical sets.

The agreement of (2.14) with experimental sideband Rabi frequency ratios measured right after the Stark shift has been nulled is within the order of 1%. The level of predictability of this ratio points to a way to measure the overlap of the red and blue tones of Raman 2 at the ion location (see section below for more information on this overlap and its significance). We control Ω_{red} and Ω_{blue} with the voltage amplitude of the RF drive of the Raman 2 AOM. For optimally aligned beams, including the Raman 2 red and blue tones, the ratio in (2.14) points to a single value for the ratio that the voltages should have[3], every time. If at some point we realize that these voltages (in the LabVIEW control software) have a different

[3]Assuming the total RF power going into the AOM is the same in all these comparisons, to account for nonlinearities of its efficiency as a function of RF power.

ratio, this indicates that the intensity ratios of the red and blue tones of Raman 2 are not as expected, most likely because of a relative misalignment between them.

2.4 Stark shift noise mechanisms

Any aspect of the light-ion interaction that is noisy leads to Stark shift noise. Intensity fluctuations lead to fluctuations of the base Rabi frequency Ω_0. Frequency fluctuations (laser or qubit frequency) cause fluctuations of the $K_{\text{single},i,j}$ and $K_{\text{cross},ij}$ factors. Finally, polarization fluctuations lead to fluctuations of the P_{ij} factors. Given realistic conditions for our lab, the intensity noise is by far the primary contributor to Stark shift noise.

2.4.1 Intensity noise

The intensity of any Raman beam at the ion location fluctuates because of global power fluctuations in that beam or its displacement relative to the ions—pointing noise—given its gaussian profile. Measurement after implementation of a power lock ("noise-eater lock") for the 355 nm optics path suggests that the power noise of a beam has a fractional standard deviation $\sigma_{\text{power}} \approx 0.6\%$ with the lock active and at the ion location[4]. As for beam pointing noise, measurements with the pointing lock active suggest fractional standard deviation σ_{pointing} between 0.5% and 2% in each direction, as a fraction of the beam waist. See section 2.4.1.1 for the pointing measurement in detail.

I will argue that power fluctuations are the main contributor to Stark shift

[4]Measurement taken at 3/4/2020.

fluctuations. This claim is founded on measurements of single-beam Stark shift noise in conjunction with power noise. At Fig. 2.6 (left plot), each point represents averaged Ramsey experiments measuring the single-beam Stark shift of Raman 2—typically minimized to a few tens of Hz, and 200 Hz for this measurement[5]. To measure it, we initialize a single ion to a state at the equator of the Bloch sphere and we turn on Raman 2 by itself, scanning this duration. We end with an analysis pulse and state readout, measuring the spin projection along X or Y. An exponentially decaying sinusoidal was fitted to its time trace and the inverse of the decay time is plotted as the decay rate. This procedure was repeated for multiple ion positions. Different ion positions see varying Stark shift coherence decay times. Using a simple Monte Carlo simulation to map those decay rates to spread of the shot-to-shot distribution of Stark shifts, we obtain a fractional Stark shift noise of about 1.5%—at least for this single-beam case.

Keeping that value in mind, let's look at how much of this Stark shift noise is explained by beam power fluctuations. From Table 2.1 we see that the shift scales with the square of the beam intensity. Importantly, if this intensity is treated like a gaussianly distributed random variable at the ion position, its square will have twice the standard deviation, and not just $\sqrt{2}$ larger. As stated earlier, this power was measured with fractional noise $\sigma_{\text{power}} \approx 0.6\%$, and therefore it is by itself causing fractional Stark shift noise $\geq 1.2\%$. In other words, it can account for 80% of the single-beam Stark shift noise. This notable contribution motivates further investigation on the theoretical modeling side, and points to the lowest hanging fruit

[5]Measurement taken at 3/2/2021.

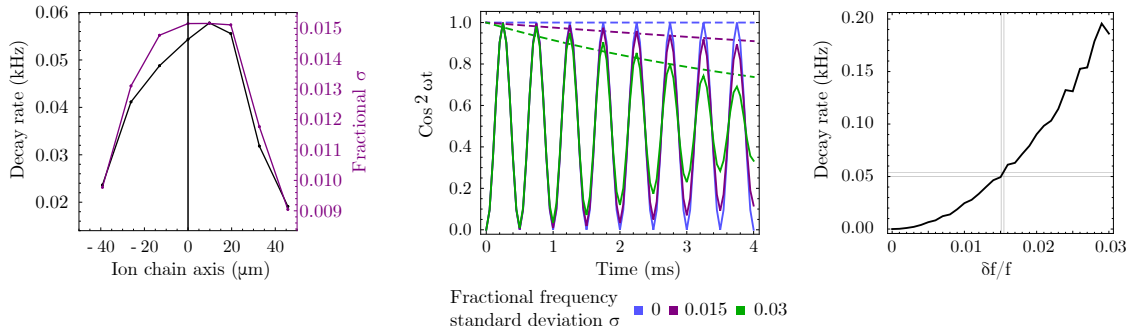

Figure 2.6: **Noise measurement and estimation for single-beam Stark shift. Left:** The decay rate of the single-beam Stark shift of Raman 2 for various ion locations along the chain axis (black). In purple, the decay rate converted to a fractional noise (standard deviation σ of the Stark shift itself. **Middle:** To convert decay rates to Stark shift noise, I average a large number of sinusoidals with a given noise standard deviation σ of their frequency (gaussianly distrubuted). Different colors are different averaged signals; more noise makes the signal decay faster. In dashed lines is the exponential decay envelope. **Right:** The dependence of the decay rate on the fractional Stark shift noise. The grid lines mark the noise case for the center of the ion chain, from the leftmost plot. The fractional Stark shift noise extracted is approximately 1.5%.

for system improvement on the experimental side. A measurement of single-beam Stark shift was chosen, because it is more straightforward to interpret compared to the cross-beam case. In general, the decay rates observed in cross-beam Stark shifts are higher than the single-beam case by a factor of 10, roughly.

For the cross-beam Stark shifts and for the sake of generally applicable results that go beyond a measurement at a specific time and lab conditions, I will use a model that allows constraining the large number of independent variables determining the shifts. There are several assumptions I will make. First, that the three relevant beams (Raman 1, Raman 2-red tone, Raman 2-blue tone) share the same amount of fractional power noise; same for their fractional pointing noise (normalized by the beam waists) in both horizontal and vertical axes. Second, that the horizontal centers of Raman 2-red and Raman 2-blue are not be perfectly over-

lapped, but they are at a distance d_{rb}, along the ion chain axis. This is motivated by our limited ability to perfectly overlap them, and a typically acceptable residual Stark shift gradient of order 10 Hz/μm along the ion chain. Third, that the ratio r_{c-i} of common-mode to independent noise between Raman 2-red and Raman 2-blue is the same for power and pointing noise. Fourth, that the Raman 1 beam's power and pointing are uncorrelated to Raman 2's homologue properties (although the distributions for each beam have the same σ for eg. power).

Note that r_{c-i} is not restricted from 0 to 1, but runs from 0 to ∞. As an example, for Raman 2, the total power noise σ_{power} is assumed to be the same for each of the red and the blue tones. Within each tone, σ_{power} is allocated to COM noise with σ_{COM} standard deviation and uncorrelated noise with σ_{unc} standard deviation, through the relation

$$\sigma_{\text{COM}} = \sigma_{\text{power}} \frac{r}{\sqrt{1 + r_{c-i}^2}} \qquad \sigma_{\text{unc}} = \sigma_{\text{power}} \frac{1}{\sqrt{1 + r_{c-i}^2}}$$

which ensures $\sigma_{\text{COM}}^2 + \sigma_{\text{unc}}^2 = \sigma_{\text{power}}^2$.

Fig. 2.7 uses these assumptions to assess the amount of Stark shift noise expected, given specific power and steering noise, as well as different ratios of COM-to-independent noise for red and blue tones. The Stark shift (vertical axis of each plot) is normalized to the cross-beam Stark shift that *one* tone, red or blue, produces; of the order of 12 kHz for typical lab operation. *A notable result is that beam pointing noise matters less that an equal amount of (fractional) power noise.* This does not come as a surprise, since the ion chain sits largely on top of the gaussian intensity

profile of the beams and even if they jitter around that position, the intensity profile does not change much at the chain location. This argument holds for horizontal pointing noise where the beam waists are large ($\approx 100\ \mu$m). As shown in 2.4.1.1, vertical pointing noise along the tight beam axis is also present but not to an extent belying the above argument. The arrow in the figure points to a realistic approximation for our lab: the pointing noise is set to 1%, the centers of the red and blue tones are assumed to be 3 μm apart, and the ratio of common-mode/uncorrelated noise for both power and pointing in the beams is assumed to be 2.

2.4.1.1 Beam pointing and ion position noise measurement

Raman beam pointing noise contributes to Stark shift noise. Once we measure the pointing noise, calculating its contribution to Stark shift noise is fairly straightforward, under minimal assumptions. We measured Raman 2's pointing noise at the ion location and presumed similar behavior for Raman 1.

The setup at Fig. 2.8 was used to look at the pointing of the red and blue tones of Raman 2 virtually *simultaneously*. The red tone by itself was switched on for 50 μs, immediately followed by the blue tone by itself for another 50 μs. This was repeated for 200 seconds. Since we do not expect to see noise faster than $1/100\ \mu\text{s} = 100$ kHz, this alternating scheme is equivalent to measuring the two tones simultaneously. To eliminate the effect of AOM and detector rise/fall times, we kept the middle third of each 50 μs pulse. The results are shown in Fig. 2.9 and their Allan variance in Fig. 2.10. For the horizontal axis, the standard deviation of

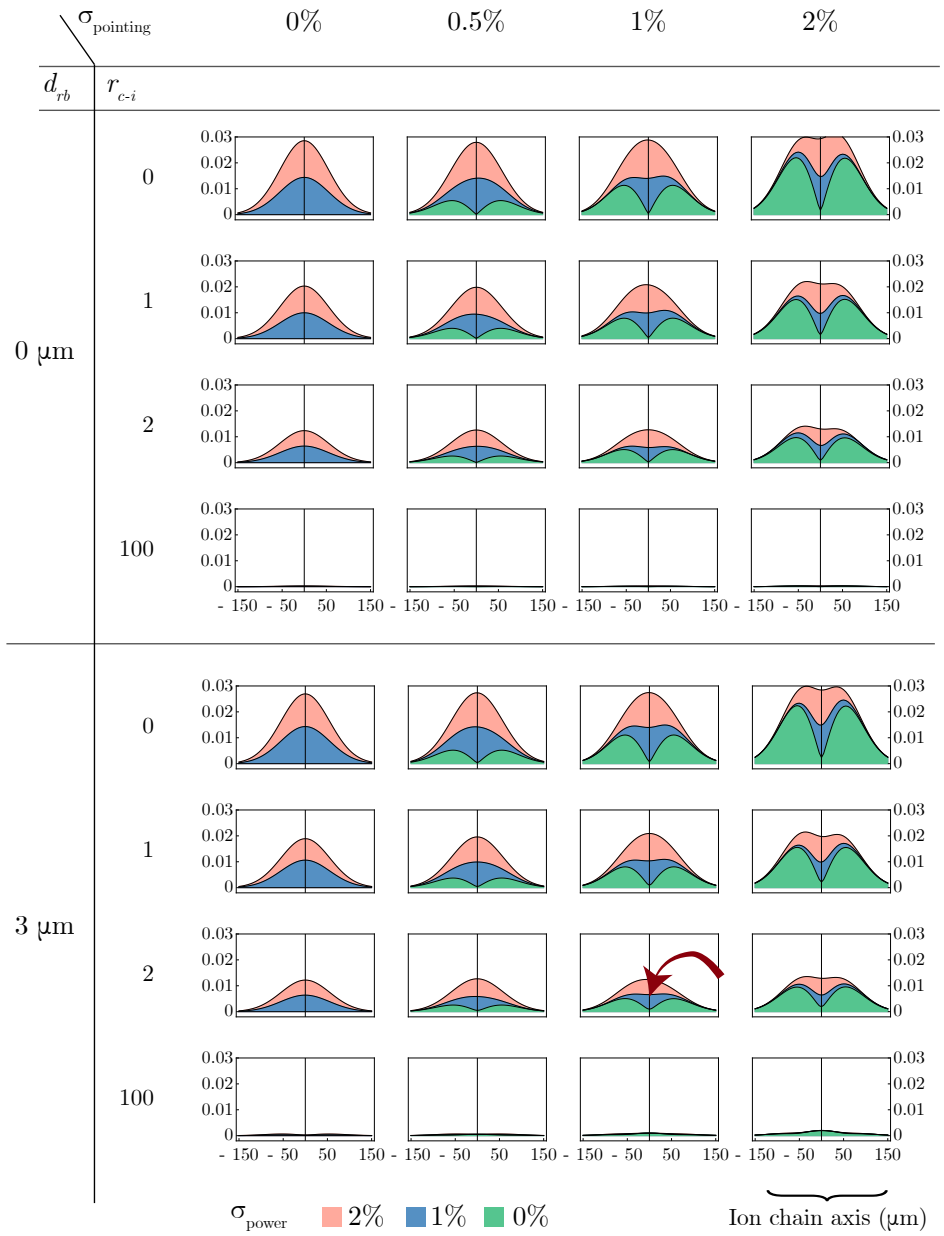

Figure 2.7: **Differential cross-beam Stark shift from both red and blue beatnotes.** Case scenarios for the noise—not magnitude—of the total cross-beam Stark shift. Each plot vertical: its fractional standard deviation divided by the cross-beam Stark shift of a single-tone. Horizontal: ion chain axis in μm. Pointing noise σ_{pointing} from 0% to 2% and power noise in each beam/tone from 0% to 2%, modeled as described in the text. A distance $d_{rb} = 0$ μm between the horizontal centers of the red and blue tones of Raman 2 implies perfect overlap, while the 3 μm value is closer to our typical tolerance. A ratio $r_{c-i} = 0$ of common mode to independent noise implies that all noise is uncorrelated between the red and blue tones, while a large value like 100 that it's mostly COM. The red arrow indicates a reasonable best-case approximation for QSim and it corresponds to total Stark shift noise of around 100 Hz in standard deviation, assuming a red or blue Stark shift of around 12 kHz.

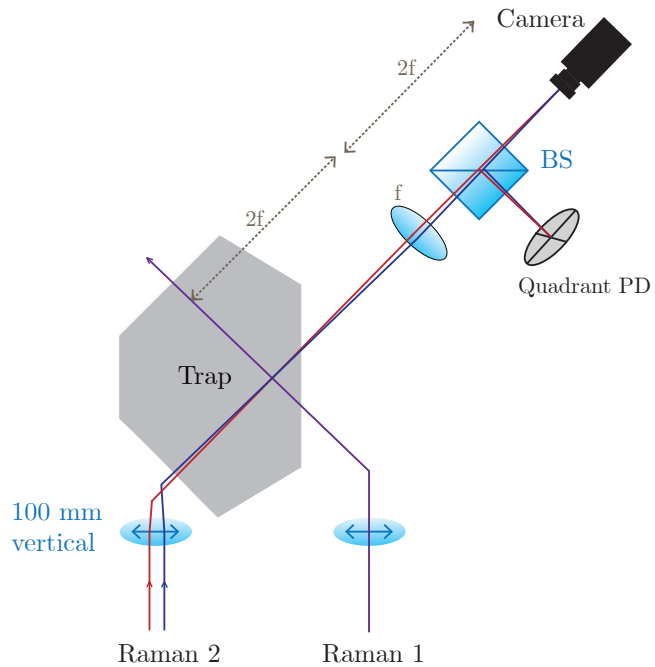

Figure 2.8: **Optics diagram for the beam pointing noise measurement.** Using a quadrant PD, we measured the steering noise of Raman 2 and its two tones, red and blue. To accurately reproduce the beam behavior at the ions, we placed a spherical converging lens of focal length f at a distance $2f$ away from the trap. Another $2f$ away from it, we collected half of the beam at the quadrant PD and the other half with a Guppy beam profiler camera. This was necessary to calibrate the quadrant PD measurement (convert its voltage reading to distance). The quadrant PD was chosen over the beam profiler for the actual measurements due to its higher sensitivity and acquisition rate. To ensure the $4f$ configuration of this measurement setup, we made sure that the red and blue tones appear to be maximally overlapping at the beam profiler, since we know they are optimally overlapped at the ion location.

the horizontal beam centers was about 0.5 μm (on top of a beam waist of about 100 μm). For the vertical axis, it was about 0.15 μm (on top of a waist of about 10 μm).

To complete the pointing noise investigation, we considered ion motion relative to the trap and the beam. A cause could for example be noise in the trap voltages. The closest quantity we could measure was motion of ions relative to the imaging system, essentially by taking a picture of a single ion every $\sim$ 10 ms (Fig. 2.11). The position fluctuations were negligible. Still, it is interesting to observe in the Allan

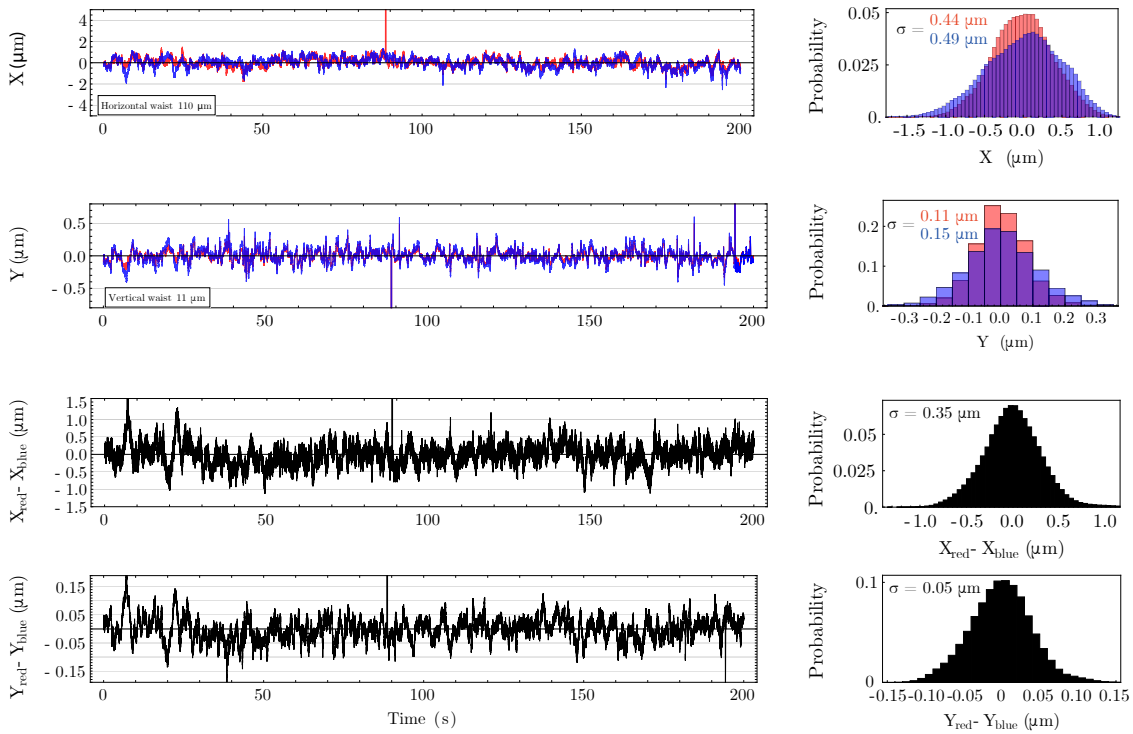

Figure 2.9: **Beam pointing noise measurement.** From top to bottom: horizontal axis (parallel to optical table and at 45° angle to the ion chain, vertical axis, differential for horizontal axis, differential for vertical axis. It seems that the noise is mostly common mode for the vertical axis, while it's more split between common and differential for the horizontal axis.

variance that drifts start occurring at timescales of 1 s or longer. Additionally, it seems that their position fluctuates slightly more horizontally than vertically. This is not related to harmonic oscillator wavefunction spreads, since these would be at least an order of magnitude smaller than the spreads observed. It is likely that along the horizontal axis, the weaker axial confinement allows for larger traversations of the ion, after a disturbance like a background collision or trap voltage noise.

In conclusion, neither beam pointing or ion position fluctuations pose as a significant source of intensity noise at the ions. Owing to the gaussian profile of the beams, the ions are sitting at a maximum with a zero spatial derivative of the

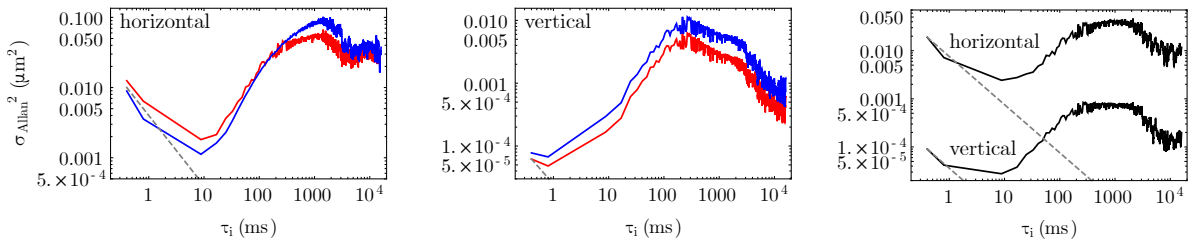

Figure 2.10: **Allan variance of the beam centers from the data of Fig. 2.9.** Left: horizontal (red and blue tones of Raman 2). Center: vertical. Right: differential.

intensity, and therefore small deviations have little effect. This is further evidenced in Fig. 2.7. A caveat in this argument is that the Raman beams have not always been measured to have an ideal gaussian profile at their focus. Burnt or dirty optics have occasionally led to the intensity profile having features at a scale much smaller than the waist, and it is possible that even pointing noise of the order ~ 1 μm could have a significant effect in that case.

2.4.2 Polarization and frequency noise

The other two fundamental sources for Stark shift noise are polarization and frequency. Fig. 2.12 presents a measurement of the polarization noise in the Raman 2 beam, measured close to the ion location. It substantiates the claim that our polarization noise is negligible as far as its effect on Stark shift is concerned.

Frequency noise is even less significant. The frequencies that could be fluctuating are the qubit frequency because of fluctuating magnetic fields, and the frequencies of the RF comb teeth. We know the former is negligible—after all, that is one of the reasons why the $^{171}\mathrm{Yb}^+$ hyperfine states form a good qubit: they are insensitive to the magnetic field at first order.

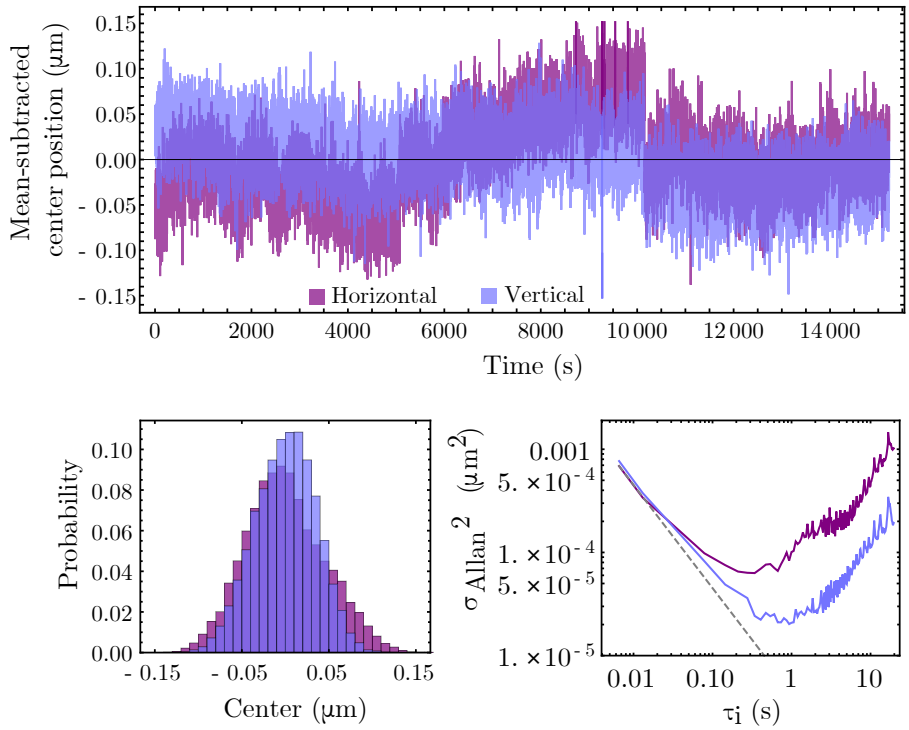

Figure 2.11: **Ion position noise.** The position of the ions relative to the imaging system shows small fluctuations with standard deviations of 0.046 μm for the horizontal (along the ion chain) direction and 0.036 μm for the vertical (perpendicular to the Raman beams plane) axis. For reference, the horizontal beam waists are about 100 μm horizontally and about 10 μm vertically.

As far as the RF comb teeth go, remember that the only thing that we are locking is our RF beatnote to the frequency of the $m = 105$th comb tooth. All the other comb teeth that are causing 4th order Stark shifts are not locked. Fig. 2.13 showed a 100 s-long measurement of the repetition rate drift, over which its fractional standard deviation was $2 \cdot 10^{-7}$, eliminating it from the list of concerns as well.

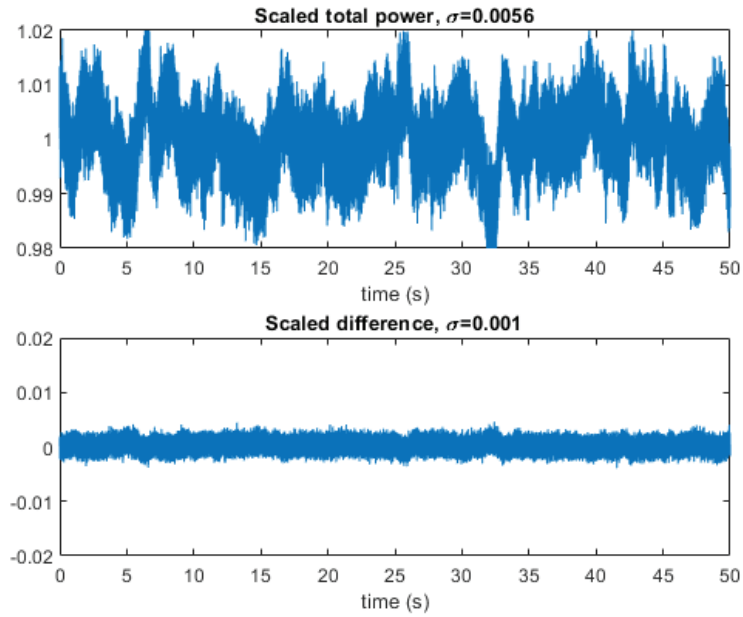

Figure 2.12: **Polarization noise in Raman 2** To measure the polarization noise in the Raman 2 beam, a PBS was placed in its path right after the trap, and the power of the transmission and reflection were measured. The upper plot is the sum of these two, proportional to the total power. The lower plot is the difference between the ports, which is the signal we care about, since the transmission is horizontally polarized and the reflection verticallly polarized. This difference is normalized to the total power from the upper plot, to scale out the effect of overall power fluctuations. After that, the fractional standard deviation is 0.001, making this a negligible effect. scaled difference

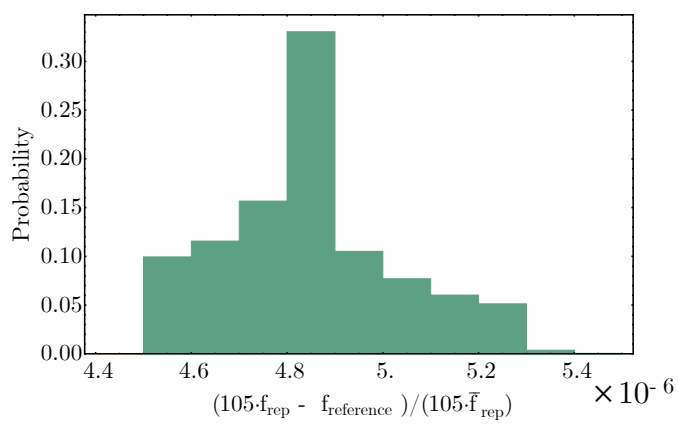

Figure 2.13: **Repetition rate noise of our pulsed Raman laser.** Using the PI beatnote lock setup, we switched the stabilization off and observed the fluctuations of the 105th tooth (mixed down) for 100 s. The fractional fluctuations of noise have a standard deviation of 10^{-7}, consistent with the manufacturer's claim for repetition rate noise of a few tens of Hz. Like polarization, this is a negligible Stark shift noise source.

Chapter 3: Other imperfections

This chapter discusses three aspects of quantum simulation with the ion trap, all relevant to departures from ideal, power-law interactions evolution, that we have had to somehow address. This implied either making technical improvements in the apparatus, or modeling error terms to match the quantum simulation results to expectations.

3.1 Phonons and loops

The Ising interaction is a derivative of the laser-ion interaction. Unfortunately, there's another player in the field: from (1.15) we have seen that additionally to the pure spin term involving χ_{ij}, there is the spin-phonon entanglement term involving $\hat{\zeta}_i$. This term resembles a product of displacement operators for each mode k, weighted by each ion's participation to each motional mode, b_{ik}. Each term $\hat{\zeta}_i$ involves [2]

$$\alpha_{ik}(t) = i\frac{\eta_{ik}\Omega_i}{2\delta_k}(e^{-i\delta_k t} - 1) \tag{3.1}$$

where the phase difference between the red and blue beat notes was taken to be $\pi/2$ (without loss of generality).

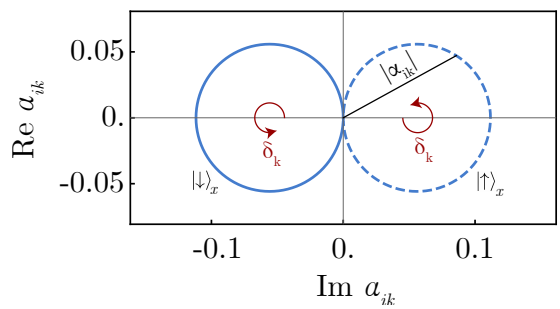

Figure 3.1: **Phonon loops for $\alpha_{ik}(t)$.** For typical experimental parameters, the radius of each loop is much smaller than 1. The distance from the origin, $|\alpha_{ik}(t)|$, is related to the square root of the average number of phonons generated by the action of $\hat{\zeta}_i$, weighted by the ion's participation to that phonon mode, b_{ik}

.

From (3.1) and we see in Fig. 3.1, $\alpha_{ik}(t)$ is complex and its time evolution is a loop at a rate of δ_k, i.e. the detuning from each motional mode k. Each normal mode of motion has its own α_{ik}-diagram, and in each diagram, each ion has its own pair of loops: the $|\uparrow\rangle_x$ and $|\downarrow\rangle_x$ parts of the spin wavefunction see a $\pm\alpha_{ik}(t)$ factor respectively (dashed and solid loops). The radius of each loop is

$$\frac{\eta_{ik}\Omega_i}{2\delta_k} = b_{ik}\frac{\eta_k\Omega_i}{2\delta_k}$$

where Ω_i is the 2-photon Rabi frequency this ion would have if light of the same intensity as the red or blue beatnote was tuned on-resonance. In the usual MS scheme as applied in QSim, this is about half of the maximum 2-photon Rabi frequency we measure, since we direct about half of the light intensity to each of the red and blue beat notes. To be more rigorous, we can write $\Omega_i = \sqrt{\Omega_i^{\text{red}}\Omega_i^{\text{blue}}}$. In practice, we directly measure the sideband Rabi frequencies for both the red and the blue sideband with one ion in the trap, namely $\eta_{\text{COM}}\Omega^{\text{red}}$ and $\eta_{\text{COM}}\Omega^{\text{blue}}$. Combining this

with knowledge of our beam sizes, we calculate the carrier Rabi frequencies from each beatnote, at each ion's location.

At times $t = 2\pi q/\delta_k$ ($q \in \mathbb{N}$), α_{ik} passes from the axes origin and is zero. At this moment, $\hat{\zeta}_i(t) = 0$ and there is no coupling to phonons. At times $t = 2\pi\frac{q+1/2}{\delta_k}$, there is maximum spin-phonon entanglement. In the typical QSim operation, where light intensity is not modulated in time as is common in quantum computation experiments, the incommensurate mode frequencies ω_k make the detunings δ_k incommensurate as well, and there is no time where all of them cross the origin at the same time[1]. We operate in the "dispersive" or "weak coupling regime" [2], where we address this issue by ensuring that the radius of the loops is small and that the spin-phonon entanglement is small in magnitude at all times:

$$\eta_{\mathrm{COM}}\Omega \gg \delta_{\mathrm{COM}}.$$

Typically, we set the detuning from the COM mode to be about 3 times larger than the largest of the red and blue sideband Rabi frequencies:

$$\delta_{\mathrm{COM}} \approx 3\eta_{\mathrm{COM}}\Omega$$

[1] One can argue that still, mathematically speaking, the total effect can be made ϵ-small for some time t_ϵ. This t_ϵ though would be impractically larger than the coherence time of our simulations.

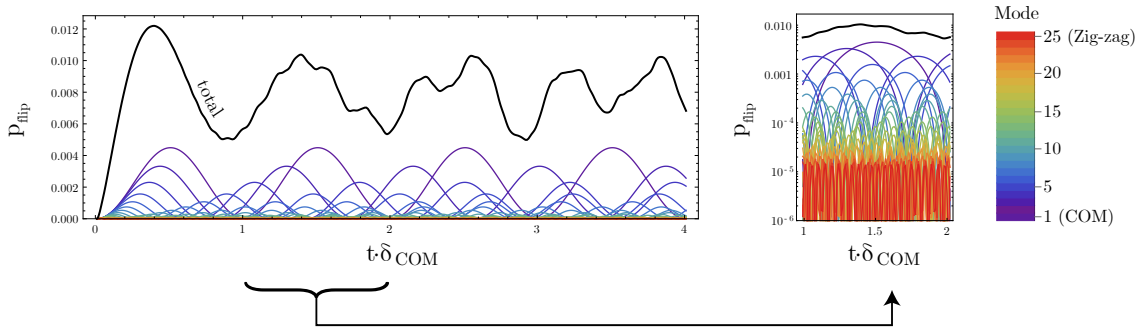

Figure 3.2: **Spin flip probability per ion, broken down by mode. Left:** the probability of a spin flip per ion p_{flip} for a chain of 25 ions and typical experimental parameters; most notably, $\delta_{\text{COM}} = 3\eta_{\text{COM}}\Omega$. The COM mode has the largest contribution. The amplitude of the sinusoidal for the k-th mode is approximately proportional to $1/\delta_k$. **Right:** To get an idea of the contribution of each mode to spin flip errors, I zoom in time and use a logarithmic vertical scale.

3.1.1 The spin-flip probabilities

Using a derivation from [35], the probability of a phonon-induced spin flip for ion i at the time of measurement t can be calculated from $\alpha_{ik}(t)$ and is

$$p_i(t) = \sum_{k=1}^{N} |\alpha_{ik}(t)|^2 \tag{3.2}$$

Fig. 3.2 shows the spin-flip probability, averaged over the chain, for typical experimental parameters. Each motional mode k contributes to a spin-flip error in the Z basis that oscillates with a frequency δ_k. The largest contribution overall is from the COM mode, since its detuning from the laser beatnote is the smallest in our typical operation—there's a higher chance we couple to it than other modes. Plotting a portion of the plot in a logarithmic scale, it is evident that the farthest detuned modes contribute negligibly to this error.

To develop an understanding of how the spin-flip probability depends on the specific ion, Fig. 3.3 shows the $\alpha_{ik}(t)$ time evolution for each ion and each mode for $N = 9$ and reasonable trap parameters. Because of the symmetric character of the motional modes around the chain center, the right half of the ion chain mirrors the left half. As the mode number increases, so does the detuning, and the loops tend to get smaller for higher modes.

It is interesting to note that merely because of the mode eigenvectors b_{ik}, the edge ions have a large participation in the tilt ($k = 2$) and $k = 3$ modes, and their spin-flip errors from each of these is larger than the COM error itself. While the same is true for, eg. ions 4 and 6 for mode $k = 8$, its large detuning suppresses its contribution to spin flips.

3.1.2 More ions, more modes, not much more error/ion

One would think that the more ions we add in the chain and the more motional modes there are to off-resonantly couple to, the higher the resulting spin-flip probability. Fortunately, physics is on our side here. As we add ions, the new transverse modes are added farther away from the laser beat note and therefore their δ_k increases; additionally, the mode spacing increases as we add more. (Fig. 3.4, left). If we also account for the inverse square dependence of spin-flip probabilities on δ_k, we see that each new mode adds less and less to the spin-flip probability for each ion, seemingly saturating.

To illustrate this statement, I am using the Poisson binomial distribution to

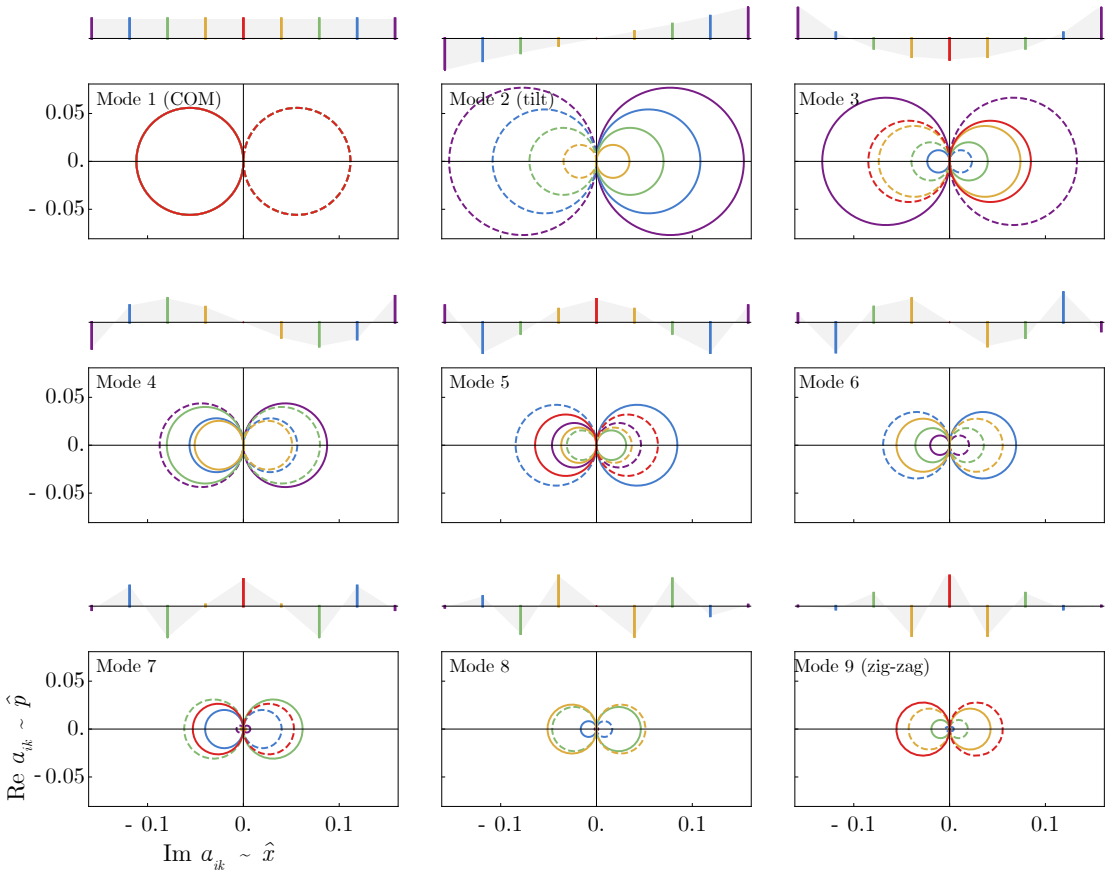

Figure 3.3: **Evolution of $\alpha_{ik}(t)$ for multiple modes, multiple ions.** All $\alpha_{ik}(t)$ plots share the same axes; similar for the mode eigenvector plots. For each plot of mode k as a whole, the radius of the loops is suppressed by $1/\delta_k^2$. Within each plot, the radius of the loops for each ion i depends on the ion's participation to the mode, b_{ik}.

calculate spin flip probabilities for an ion chain, where each ion is associated with a different flip probability (Fig. 3.4).

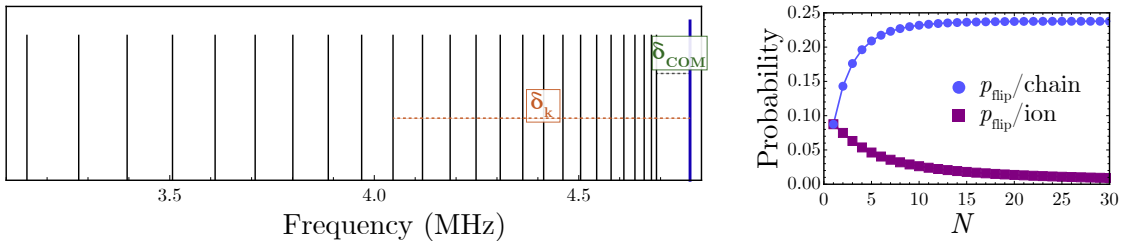

Figure 3.4: **Scaling of phonon-induced spin flips with number of ions. Left:** Frequency spectrum of transverse motional modes at the blue side of the carrier. The laser beat note (blue) is detuned by δ_{COM} from the COM mode. As we add more ions, the transverse modes add farther and farther towards the left, away from the beatnote, therefore $1/\delta_k^2$ decreases dramatically. **Right:** As we add more ions, the probability of one or more spin flips anywhere in the chain saturates because of the $1/\delta_k^2$ dependence. Correspondingly, the spin flip probability *per ion* decreases.

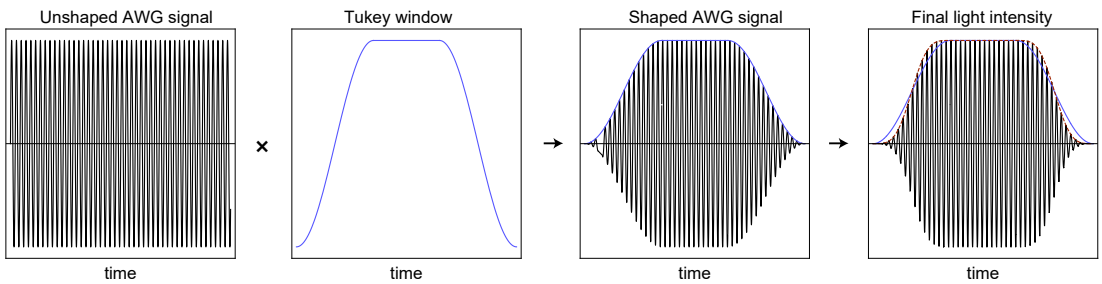

Figure 3.5: **Pulse shaping**. We multiply the original unshaped AWG signal with a smooth-edge window like the Tukey one (blue line), and get the shaped AWG signal thatis sent to the amplifiers and ultimately to the AOM. The intensity of the light that the ions see acquires a slightly modified envelope (dashed red) because of nonlinearities in the process (see next section and Fig. 3.7). Note that for this figure, it is assumed that the flat part of the pulse, i.e. its maximum value, is at the saturation power of the AOM, around 2 W for our Raman AOMs. The exact way the light intensity ramps to its maximum value depends on what this value is, because of the AOM nonlinearity of the light intensity as a function of the RF input power—see (3.5).

3.2 Finite pulse width

Single-frequency signals belong to textbooks. Signals that last finite time contain a range of frequencies, even if they are made up from a single monochromatic sinusoidal. These other frequencies may excite transitions whose frequency is close to the laser beatnote frequency, thus introducing undesired Hamiltonian terms. The first suspect is such undesirable coupling to motional modes, laying δ away from the laser beatnote in the MS scheme. Secondarily, the same beatnote could couple to the carrier transition, μ away, although the spectral component there is strongly suppressed.

To quantify the frequency decomposition of our signals, I use the truncated power spectral density (PSD) $S_T(\omega)$, based on the Fourier transform $F_T(\omega)$ of a

signal $f(t)$ which lasts for time T:

$$S_T(\omega) = \frac{1}{T}\langle |F_T(\omega)|^2 \rangle \qquad F_T(\omega) = \int_0^T e^{-i\omega t} f(t)\mathrm{d}t \tag{3.3}$$

The ensemble average symbol $\langle \ldots \rangle$ is especially meaningful when the signal $f(t)$ is noisy; in our case it is tied to an AWG-calculated signal and we can safely disregard the averaging. In the simplest case, the RF beatnote contains 2 frequencies, f_{red} and f_{blue}, and lasts for the duration of the interaction, τ_{quench}. The red and blue frequencies are 2μ away from each other and spectral components trailing off one of them are severely attenuated at the location of the other. Therefore, when considering coupling to the blue sidebands, we only need take into account the spectral structure of the blue beatnote, and vice versa for the red side.

To visualize, let's look at the blue side of the spectrum in Fig. 3.6. The PSD of the pulse peaks δ away from the blue sideband f_{blue}, but is not zero at f_{blue} itself. Therefore, the blue sideband will be weakly driven. To estimate how much, I define r_{unsh} using the PSD of the unshaped signal:

$$r_{\text{unsh}} \equiv \frac{S_T^{\text{unsh}}(\omega_{\text{COM}})}{S_T^{\text{unsh}}(\mu)}$$

For $\delta = 60$ kHz and $\tau_{\text{quench}} = 40\mu$s, $r_{\text{unsh}} = 0.019$ (Fig. 3.6b left). If the beatnote is on resonance with the COM mode ($\delta = 0$ kHz), the BSB Rabi frequency is 20 kHz. Therefore, for $\delta = 60$ kHz, we expect it to be excited with a $0.019 \cdot 20 = 0.38$ kHz Rabi frequency, assuming no pulse shaping. Note that this is a rather unfavorable

choice for τ_{quench}: relatively short, which means that its spectrum will spread more around the beatnote frequency).

To mitigate this effect, we shape our pulses using the Tukey window $w(t)$ with a ramp time typically set to $t_{\text{ramp}} = 10\mu$s (Fig. 3.5):

$$w(t) = \begin{cases} 0 & , t < 0 \text{ or } t \geq \tau_{\text{quench}} \\ \sin^2 \frac{\pi t}{2t_{\text{ramp}}} & , t \in [0, t_{\text{ramp}}) \\ 1 & , t \in [t_{\text{ramp}}, \tau_{\text{quench}} - t_{\text{ramp}}) \\ \sin^2[\pi/2(2 + (t - \tau_{\text{quench}}))] & , t \in [\tau_{\text{quench}} - t_{\text{ramp}}, \tau_{\text{quench}}) \end{cases} \tag{3.4}$$

This allows us to suppress the PSD of the laser beatnote at the COM mode by a factor of r_{impr}, defined as

$$r_{\text{impr}} \equiv \frac{S_T^{\text{shaped}}(\omega_{\text{COM}})}{S_T^{\text{unsh}}(\omega_{\text{COM}})}$$

Multiplying $r_{\text{unsh}} \cdot r_{\text{impr}}$ gives the coupling to the COM mode with shaped pulses, and in general is brought down to $\lesssim 10$ Hz, assuming SB Rabi frequencies around 20 kHz.

Interestingly, a non-negligible process to take into account when calculating S_T^{shaped} is the nonlinearity of the relative AOM efficiency ϵ_{rel} as a function of the RF drive power[2]:

$$\epsilon_{\text{rel}} = \sin^2 \left(\frac{\pi}{2} \sqrt{\frac{P_{\text{in}}}{P_{\text{sat}}}} \right) \tag{3.5}$$

where P_{in} is the driving RF power and P_{sat} the saturation RF power. To illustrate the effect of this nonlinearity, the PSD of the "naive" shaped signal, i.e. a pure Tukey-shaped signal, is also shown in Fig. 3.6.

In principle, one could aim to tune τ_{quench} such that the COM mode frequency

[2]ISOMET application note: `http://www.isomet.com/App-Manual_pdf/AOFRQSHF.pdf`

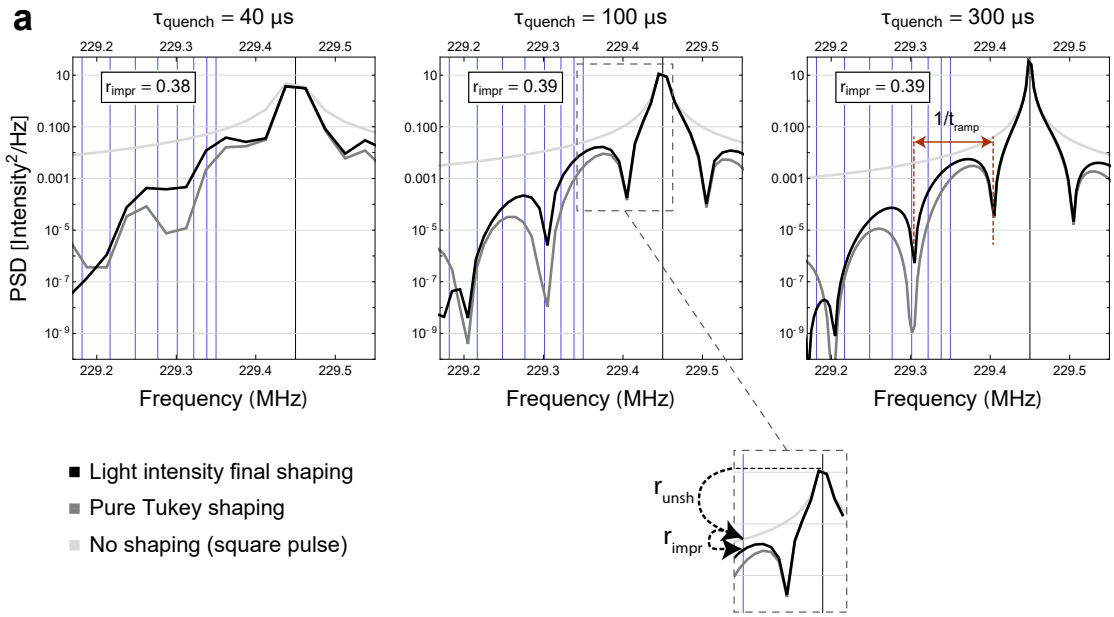

b

δ = 60 kHz			
τ_{quench}	40	100	300
r_{unsh}	0.0190	0.0029	0.0003
r_{impr}	0.2100	0.1200	0.1300
r_{tot}	0.00402	0.00035	0.00004

δ = 100 kHz			
τ_{quench}	40	100	300
r_{unsh}	0.0067	0.0010	0.0001
r_{impr}	0.3768	0.3880	0.3870
r_{tot}	0.00251	0.00039	0.00004

Figure 3.6: **Power spectral densities for pulses of 3 different durations** τ_{quench}. **a.** The light intensity PSD for the shaped pulse (black) is more tightly concentrated around the beatnote frequency compared to the unshaped light pulse (light gray). The dark gray PSD corresponds to the pure Tukey shaping, without accounting for the nonlinearities introduced by the AOM and its efficiency dependence on the driving RF power. It is included in the plot to show that in principle it is not an insignificant detail. The improvement ratio r_{impr} does not significantly change with the quench duration. The trap (blue) motional frequencies for a typical configuration are shown in blue (COM transverse frequency: 4.85 MHz, COM axial frequency: 0.335 MHz). The detuning δ of the beatnote from the COM mode in this case was 100 kHz.
Zoom-in inset: To quantitatively characterize the frequency spillover to the COM mode with the square pulses and the improvement from the pulse shaping, I define r_{unsh} and r_{impr} respectively. **b.** Relevant rations for 2 motional detunings, δ, 60 and 100 kHz.

lies centered inside the closest "valley" to the central peak at Fig. 3.6a. Since the valley center is approximately $1/(2t_{\text{ramp}})$ away from the peak, they would have to set $1/(2t_{\text{ramp}}) \approx \delta$. For $\delta = 100$ kHz, we would get $t_{\text{ramp}} \approx 5\mu$s.

3.2.1 The pulse shaping effect on the strength of the simulated Hamiltonian

Alas, improvements often come with trade-offs. In the case of pulse shaping, the trade-off is a reduction in the strength of the simulated interactions and transverse fields. Since the light intensity is not at its maximum value during the ramp-up and ramp-down times, the Hamiltonian terms involving spin-spin interactions, as well as transverse fields along X or Y, will be reduced. The dependence of these two terms on the light intensity is different, so their reduction differs as well. To complicate matters more, the Tukey window is applied to the waveform calculated at the AWG; to calculate its effect on the light intensity one must account for the intermediate stages of amplification and AOM diffraction (Fig. 3.7a). Remembering that the Tukey window $w(t)$ is only applied to the Raman 2 beam and Raman 1 is unshaped, we eventually get:

$$I_2(t) = I_2^{\max} \cdot \sin^2\left[\frac{\pi}{2} v_{RF} w(t)\right] \tag{3.6}$$

where I_2 is the intensity of Raman 2 and we used the relation for the AOM efficiency (3.5). The maximum intensity $I_2^{\max}$ depends on the input light intensity of Raman 2 and the AOM efficiency for that RF frequency, measured at saturation RF power.

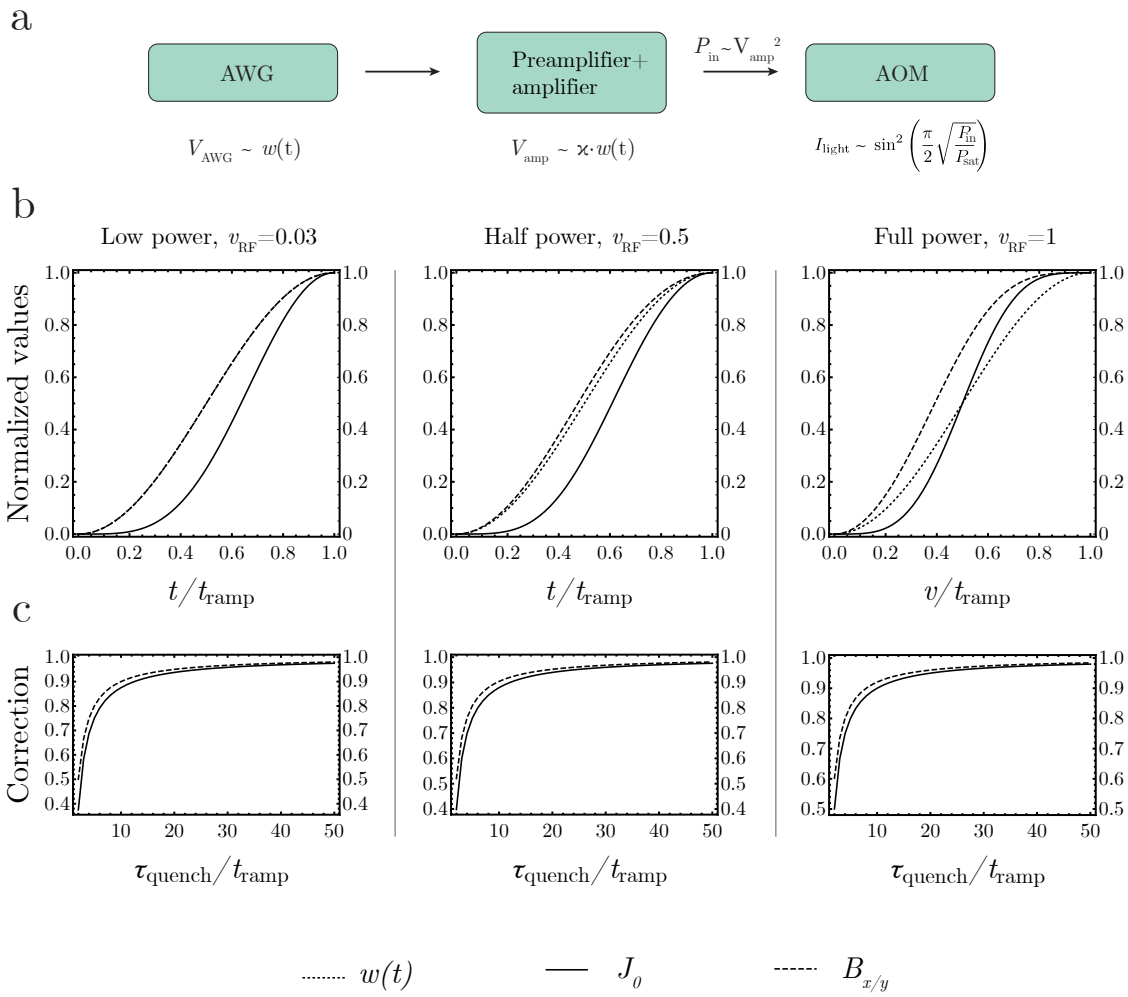

Figure 3.7: **Calculating the effect of pulse shaping on the simulation Hamiltonian. a.** First, the AWG waveform with the Tukey envelope $w(t)$ is amplified with a pre-amplifier and an amplifier; it is assumed that the output $\kappa w(t)$ is linear to the input $w(t)$. The AOM then transforms the input RF power P_{in} into diffracted light power according to (3.5). P_{sat} is the optimal (maximum) RF power we can drive the Raman AOM, which for us is 2 Watts. **b.** Ramp profiles for the Tukey window $w(t)$ (dotted line), the transverse fields along X or Y (dashed line) and the interaction $J_0(t)$ (solid line), normalized to their maximal values for each specific v_{RF}. Each of those is shown for three different RF drive powers v_{RF}, corresponding to a small drive of $v_{RF} = 0.03$, a drive at half the saturation power roughly equal to the case for each of the red and bue sidebands in the MS scheme, $v_{RF} = 0.5$, and the maximum drive $v_{RF} = 1$. **c.** To calculate the effective strength of each of the 2 terms (interaction and transverse fields), we need to integrate the pulse time profile. The longer the pulse, the better estimate the maximum value of the term is, since the ramp duration t_{ramp} becomes proportionally smaller in comparison to the total duration τ_{quench}.

The ratio v_{RF} is the voltage amplitude applied to the AOM RF input, divided by the voltage amplitude corresponding to saturation power. I use v_{RF}, because it is the parameter we directly control with our LabVIEW control software. For example, when driving the AOM with full RF power, $v_{RF} = 1$ and at the end of the ramp we have $I_2(t = t_{\text{ramp}}) = I_2^{\max}$. For low RF drive of $v_{RF} = 0.03$, corresponding to a software input parameter of 1000[3], we get $I_2(t = t_{\text{ramp}}) \approx 0.002 I_2^{\max}$.

The transverse fields $B_{x/y}$ along X or Y are proportional to the Raman Rabi frequency Ω, while the interactions strength $J_0 \sim \Omega^2$. Remembering that $\Omega \sim \sqrt{I_1 I_2}$, the interaction strength scales as $J_0 \sim I_1 I_2$, and we end up with

$$B_{x/y}(t) = B_{x/y}^{\max} \cdot \sin\left[\frac{\pi}{2} v_{RF} w(t)\right] \tag{3.7}$$

$$J_0(t) = J_0^{\max} \cdot \sin^2\left[\frac{\pi}{2} v_{RF} w(t)\right] \tag{3.8}$$

We substitute the time-dependent Hamiltonian terms above with static ones, derived from the maximum values $B_{x/y}^{\max}$ and $J_0^{\max}$ rescaled by a factor smaller than 1 to give an effective value of the strength of that term for a given quench duration. For example, the corrected X or Y field would be

$$\frac{1}{\tau_{\text{quench}}} \int_0^{\tau_{\text{quench}}} B_{x/y}(t) \mathrm{d}t = C(v_{RF}, \tau_{\text{quench}}/t_{\text{ramp}}) \cdot B_{x/y}^{\max} \tag{3.9}$$

The correction factors $C(v_{RF}, \tau_{\text{quench}}/t_{\text{ramp}})$ for three values of v_{RF} and for a range of quench durations are plotted in Fig. 3.7c.

[3] out of 32767, which is the maximal value

3.3 The power law approximation for the interactions

The power law approximation is the statement

$$J_{ij} = \frac{J_0}{|i-j|^\alpha} \tag{3.10}$$

modelling how the simulated spin-spin interactions depend on the distance between each pair of spins $\{i, j\}$. This is a heuristic fit result rather than an actual relation between the elements of the J matrix. In reality, neither are the couplings invariant along the chain, nor are they exactly described by a power law. In practice, the power law is a good approximation for small numbers of ions ($N \lesssim 20$).

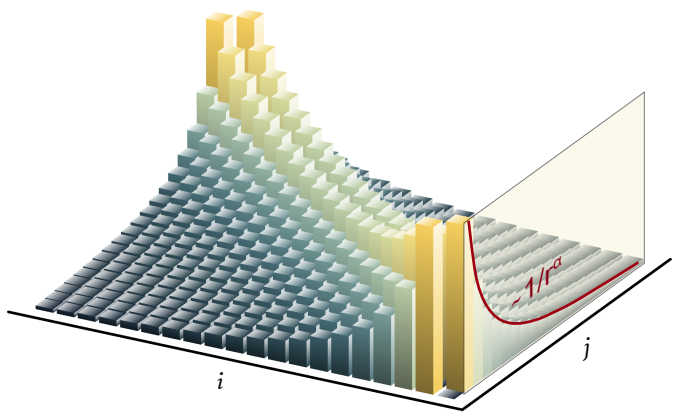

Figure 3.8: **Sample** J_{ij}. The position of each column represents a pair of spins $\{i, j\}$ and its height (also color-coded) represents the strength of that coupling. The shaded panel shows how the spin at one end of the chain ($i = 1$) couples to the rest ($j = 2, 3, \ldots, N$). The coupling is approximated by a power law with interaction range α.

In addition, the J_{ij} couplings are not uniform along the chain (parallel to the

"ridge"), as illustrated in Fig. 3.8:

$$J_{i,i+k} \neq J_{j,j+k} \qquad \text{for } i \neq j.$$

The profile of J_{ij} is dictated by the normal modes by means of their eigenvectors. By tuning the bandwidth of the transverse modes, we can make the interaction matrix more homogeneous. For an axial frequency of ≈ 0.2 MHz we achieve optimal homogeneity, as shown in Fig. 3.9, given the rest of our parameters. Of course, ion chains of only up to a certain length will be linearly arranged in this axial potential and adding more ions will cause the chain to "buckle", so in practice we are not able to accommodate the homogeneity restriction.

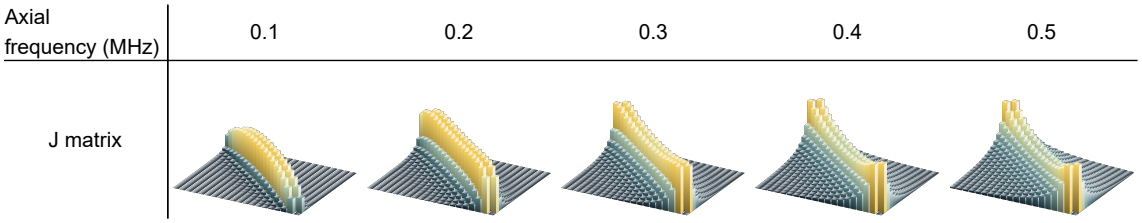

Figure 3.9: **Interactions inhomogeneity along the chain**. As the axial frequency is increased, the profile of the nearest-neighbor interactions changes from convex to concave. These matrices were calculated for $N = 20$ ions and transverse COM frequency 4.7 MHz.

Chapter 4: Prethermal discrete time crystals

4.1 Overview

This chapter features a presentation of the intriguing phenomenon of prethermal time crystals (PTC), whose experimental investigation was possible with our trapped-ion quantum simulator over recent years. My goal is not to provide a complete theoretical description, but to offer an intuitive picture of the physics involved and the experimental observations.

4.2 Expanding the definition of phases of matter

Time crystals are a phase of matter—a non-equilibrium one. The typical framework of phases of matter presumes matter in equilibrium. To expand this framework to matter that is driven, the notion of *crypto-equilibrium* [36] comes in handy. As illustrated in Fig. 4.1, a system that is time-dependent may look as if it is in equilibrium to an observer that is also time-dependent in a carefully chosen way. Therefore, there are time-dependent Hamiltonians yielding time-periodic dynamics (perhaps after some short initial transient period) that this rotating observer would validly claim to be a stationary state, enabling them to dub what they see a phase

of matter.

Additionally, a *prethermal* phase of matter is a non-trivial state of the system that appears to be stationary for very long times, even though it eventually gives way to the trivial, infinite temperature state[1]. A condition for the existence of prethermal phases ("prethermalization") is that the dynamics host a time interval during which the system's entropy is virtually fixed, despite the time-dependent Hamiltonian. For the experiment I will discuss, the prethermal regime can be made arbitrarily long.

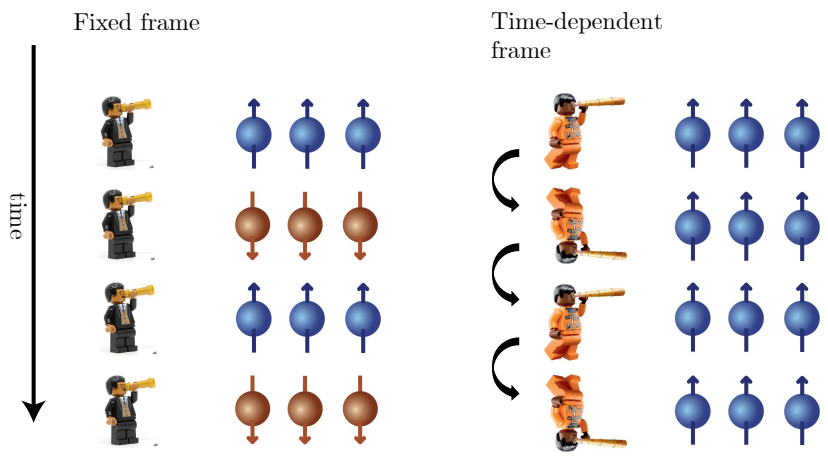

Figure 4.1: **Crypto-equilibrium.** If there is a rotating reference frame where the system does not appear to be changing, then this observer may see a stationary state and they can talk about a phase of matter.

[1]The infinite temperature state is a superposition of all Hamiltonian eigenstates with equal probability. Since the dynamics we will be considering are unitary and we begin from a pure state, this "infinite temperature" state is still a pure state. However, all observables consisting of local terms have the same value as if the system was in thermal equilibrium with an infinite temperature reservoir—hence the liberty to talk about temperature in a closed, non-equilibrium system setting.

4.3 Experimental procedure

Our system consists of 25 $^{171}\mathrm{Yb}^+$ ions, each of which represents a spin-1/2. The initial states are spin product states, most of which have all spins pointing along the positive or the negative X axis of the Bloch sphere. In order to prepare those states, we took advantage of a tightly focused laser beam to flip individual spins relative to their neighbors [29] (Fig. 4.2).

The Floquet drive consists of multiple repetitions of the following building block, itself including two distinct evolutions U_1 and U_2:

$$\begin{aligned} U_1 &= \exp\left[i\frac{\pi}{2}\sum_i^N \sigma_i^y\right] \\ U_2 &= \exp\left[iT\left(\sum_{i<j}^N J_{ij}\sigma_i^x\sigma_j^x + B_y\sum_{i=1}^N \sigma_i^y + B_z\sum_{i=1}^N \sigma_i^z\right)\right], \end{aligned} \tag{4.1}$$

where σ_i^q is the q-th component of the Pauli operator for the i-th spin, and I set $\hbar = 1$. $J_{ij} > 0$ is the long-range coupling with average nearest-neighbor interaction strength $J_0 = 2\pi \cdot (0.33 \pm 0.02)$ kHz, while $B_y = 2\pi \cdot 0.5$ kHz and $B_z = 2\pi \cdot 0.2$ kHz are global effective magnetic fields. The Floquet unitary $U_F = U_2 U_1$ implements the dynamics over a period of the drive and has frequency $\omega \equiv 2\pi/T$.

4.4 The prethermal regime

Prethermalization in our experiment relies on the inability of the system to exchange energy with the drive U_F. To intuitively understand this inability, note

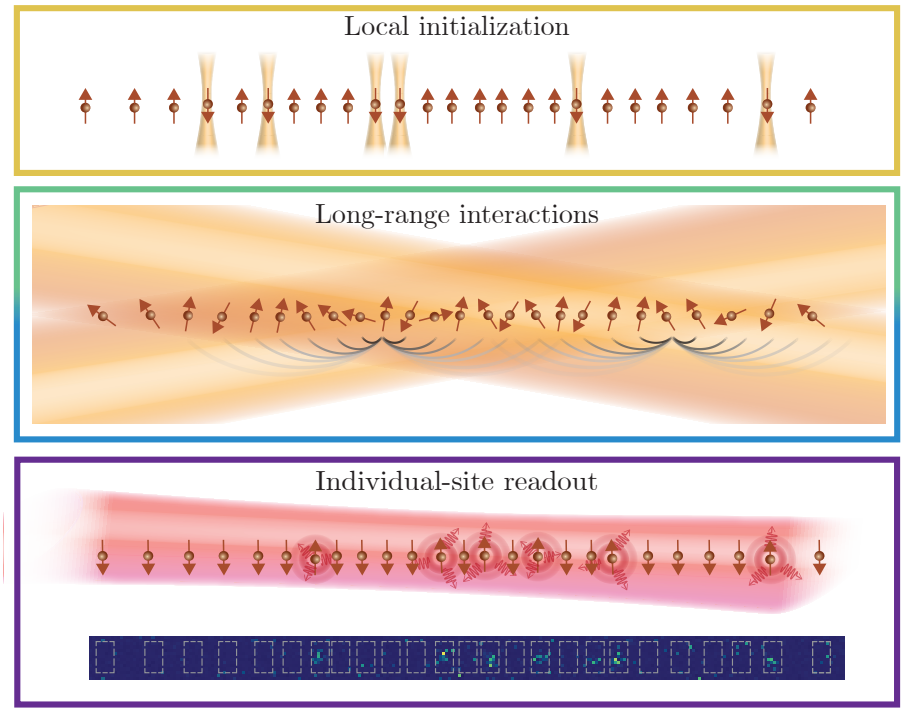

Figure 4.2: **Experimental procedure.** First, the spins are initialized using optical pumping and Raman sideband cooling, followed by single-spin rotations with the aid of the individual addressing beam shown here. Then (center panel), the Floquet unitary includes long-range interactions generated with two Raman beams illuminating the ion chain. Finally (bottom panel), the spins are projectively measured using state-dependent fluorescence. Spins that are projected to the $|\uparrow\rangle_z$ state by the detection beam, shown in red, will fluoresce, while spins projected to the $|\downarrow\rangle_z$ state will not. The fluorescence photons are collected by the imaging system and focused on the sensor of an EMCCD camera (Andor iXon Ultra)—see Section 1.1.2 for details. The presence of a few bright pixels in the region of interest (ROI) for each ion denotes an $|\uparrow\rangle_z$ state. For more details, see Appendix B. Figure adapted from [37].

the large difference between two energy scales. The first relates to the drive and is characterized by its frequency ω. The second relates to a local energy scale J_{loc} and is associated with the energy required to flip a single spin. When $\omega \gg J_{loc}$, the system cannot efficiently absorb or emit energy to the drive and undergo such rearrangements [38, 39], and therefore energy is conserved. The ratio ω/J_{loc} in the experiments described ranges from 12 to 67.

The onset of the prethermal regime is marked by a timescale noted as τ_{pre}. By τ_{pre}, the drive has vanished inhomogeneities of the initial state, and local observables are described by an equilibrium ensemble of inverse temperature β. For example, the probability of observing an eigenstate $|i\rangle$ of energy E_i is proportional to $e^{-\beta E_i}$. To experimentally observe this local equilibration stage, we prepare the chain pointing towards X with the two central spins at a 90° angle, pointing along Z. We then apply the Floquet drive U_F and note the time when the system becomes homogeneous with regard to X-magnetization (see Fig. 4.3). It is important to point out that τ_{pre} comes before decoherence extinguishes any meaningful signatures. The equilibration of the system to a quasi-thermal ensemble of the effective Hamiltonian H_{eff} (see (4.2)) ensures that around τ_{pre} and later, the observed dynamics truly correspond to the prethermal regime and are not transients mistaken for prethermal dynamics.

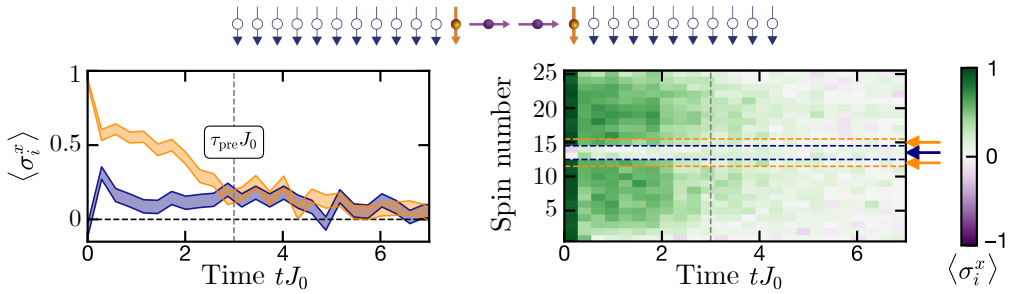

Figure 4.3: **Local equilibration and the onset of the prethermal regime. Left:** The X-magnetization of the two central spins every other Floquet period (blue) starts from zero, since these spins were initialized along Z. Their nearest neighbors (orange) start from $+X$. Around $\tau_{\text{pre}} \approx 3/J_0$, these two signals merge, meaning that the spin chain has become homogeneous in terms of spin orientation, and all spins flip in unison for the subsequent Floquet periods. This homogenization marks local equilibration; the system is hence described by an equilibrium ensemble of a specific inverse temperature β. Importantly, the value the two signals converge to is statisically different from zero, showing that decoherence is not (the only) factor explaining this convergence. **Right:** The single-spin X-magnetization data where the left plot came from. Figure adapted from [37].

The prethermal window lasts approximately until a timescale denoted as τ^* (see next paragraph). From τ_{pre} to τ^*, the stroboscopic dynamics of the system (every other period) are well-approximated by an effective prethermal Hamiltonian, which to lowest order in $1/\omega$ is given by [39]:

$$H_{\text{eff}} = \sum_{i<j}^{N} J_{ij}\sigma_i^x\sigma_j^x + B_y \sum_{i=1}^{N} \sigma_i^y. \tag{4.2}$$

The end of the prethermal regime is characterized by the timescale associated with the frequency-dependent Floquet heating, τ^*. To experimentally investigate τ^*, we measure the dynamics of the prethermal energy density $\langle H_{\text{eff}}\rangle/(NJ_0)$ for two different initial states on opposite ends of the many-body spectrum of H_{eff}: a low-energy Néel state (Fig. 4.4, left) and a high-energy polarized state (Fig. 4.4, center) along the Ising interaction axis. In both cases, increasing the driving frequency suppresses the heating rate, as expected (Fig. 4.4, right). For sufficiently large frequencies, we observe a plateau in τ^*, suggesting the presence of external noise. In particular, the origin of this plateau is consistent with experimental fluctuations of the light intensity, leading to noisy AC Stark shifts on the qubit transition. For more details on this noise, see section 4.6 and Chapter 2.

4.5 Crossing the boundary: above and below the critical energy

After establishing the existence of a prethermal time window when thermalization is slowed down for fast drive frequencies, let's look at interesting physics that is happening during that interval. At τ_{pre}, the system is described by an equilibrium

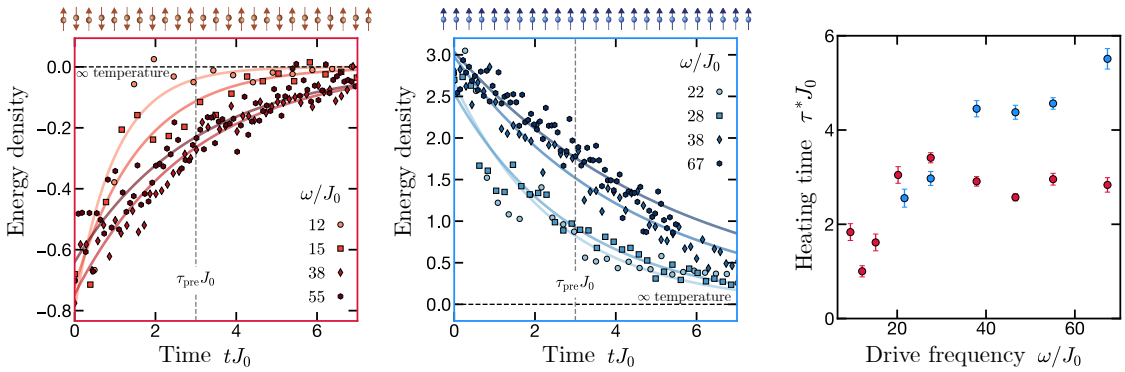

Figure 4.4: **The prethermal regime. Left & center:** Energy density evolution of a low-energy Néel state and a high-energy polarized state, respectively. Both cases show Floquet heating toward the infinite temperature value of zero magnetization, albeit from opposite sides of the many-body spectrum. In addition, in both cases, faster drive frequencies ω suppress the heating rate. Statistical error bars are of similar size as the point markers. **Right:** The heating time τ^* for the Néel (red) and polarized (blue) states increases with frequency. τ^* is extracted by fitting the dynamics of the energy density to $\sim e^{-t/\tau^*}$; exponential fits are shown as solid curves at the left and center plots. At high frequencies this behavior saturates owing to the presence of external noise. Error bars for the heating time correspond to fit errors. Figure adapted from [37].

ensemble of $H_{\rm eff}$, an antiferromagnetic Ising Hamiltonian. For this Hamiltonian, states with energy close to the most excited state have spins mostly aligned parallel to each other and along X (Fig. 4.5). This can happen with states where spins are mostly aligned towards $-X$, or where spins are mostly aligned towards $+X$. Therefore, a symmetry-breaking phase exists at the top of the spectrum at high energies (Fig. 4.6C), where spins align either towards $-X$ or $+X$.

A short interlude may be needed here for readers overly familiar with ferromagnetic interactions and associated phase transitions. Long-ranged ferromagnetic Ising interactions stabilize a low-temperature one-dimensional ferromagnetic phase but antiferromagnetic long-ranged Ising interactions do not stabilize a low-temperature one-dimensional antiferromagnetic phase. Thus, the antiferromagnetic interactions

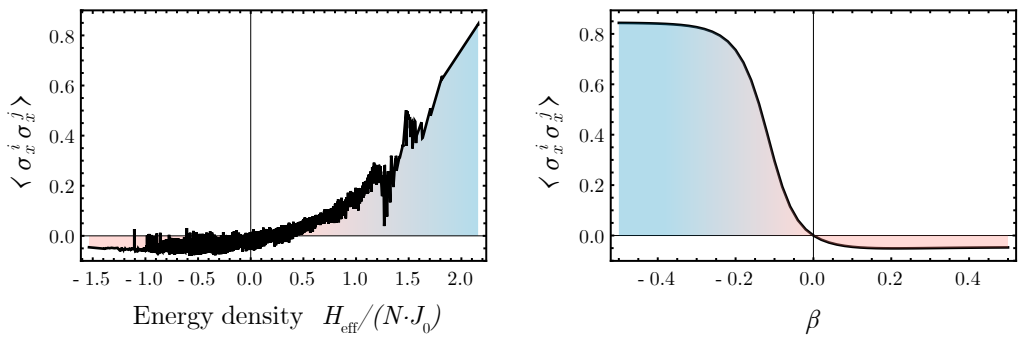

Figure 4.5: **FM regime for 14 spins.** To provide some intuition about temperature and the FM (aka, symmetry-breaking) regime in our model, I am showing calculations for the average X-magnetization spatial correlator $\langle \sigma_x^i \sigma_x^j \rangle$ for spins i and j. All experimental parameters were kept the same as with the 25-ion experiment for this calculation; the slightly different value for J_0 was compensated by adjusting B_y so that their ratio is the same. **Left:** Each of the 2^{14} energy eigenstates has a different $\langle \sigma_x^i \sigma_x^j \rangle$. Labeling eigenstates with their energy density normalized to J_0 for the horizontal axis, note that the higher the energy density, the closer to 1 the correlator (blue shading). A value of 1 signifies perfect ferromagnetism, i.e. spins aligned to each other. On the contrary, towards low energy densities on the left, the correlator is almost zero—this is a paramagnetic regime (red shading). **Right:** Using the plot on the left, we can calculate the value of the correlator for thermal ensembles of various inverse temperatures β. The ferromagnetic regime lives at negative β values, i.e. at negative temperature. Normally and in the proper meaning of temperature, this would be an inaccessible regime, since there is no thermal bath with a negative temperature. However, in our closed quantum system they are defined.

in our system do not support a symmetry-breaking phase at low energy densities but rather at the top of the spectrum at high densities (Fig. 4.6C), which is the low-energy-density regime of the ferromagnetic Hamiltonian $-H_{\text{eff}}$.

Continuing with the high energy case, we can see how breaking the symmetry $G \equiv \prod_{i=1}^{N} \sigma_i^y = U_1$ at the top of the spectrum gives rise to a time crystal. Remember that energy is conserved until τ^* and therefore the energy of the initial state we prepare will be the average energy of the ensemble at τ_{pre} (Fig. 4.6C). That ensemble though features two symmetry-breaking sectors: the "$-X$" and the "$+X$". At τ_{pre} it reaches one of these. Now, for a single period of evolution, the exact Floquet

dynamics are approximately generated by acting with G, which is a π-pulse around Y, followed by evolving under H_{eff} for time T [40, 39]. If the system has a net X magnetization at τ_{pre}, then this evolution will flip it every period. The $2T$-periodic resulting signal is a signature of a time crystal.

If, on the other hand, the initial state energy was below a critical value, the prethermal ensemble at τ_{pre} will be in the paramagnetic regime and it will have zero net X-magnetization. In other words, it will preserve G as a symmetry. The action of a single Floquet evolution will not have any effect.

We investigate these two regimes by measuring the auto-correlation of the magnetization:

$$M(t) = \frac{1}{N}\sum_{i=1}^{N}\langle\sigma_i^x(t)\rangle\langle\sigma_i^x(0)\rangle. \tag{4.3}$$

Starting with a low-energy-density Néel state (Fig. 4.7A), we observe that $M(t)$ quickly decays to zero at τ_{pre}, in agreement with the expectation that the system equilibrates to the symmetry-unbroken, paramagnetic phase. This behavior is frequency-independent, in direct contrast to the Floquet dynamics of the energy density (Fig. 4.4A). This contrast highlights an essential point: although τ^* can always be extended by increasing the driving frequency, if the system lives in the trivial Floquet phase, no order will survive beyond τ_{pre}.

The Floquet dynamics starting from the polarized state are notably different (Fig. 4.7B). First, $M(t)$ exhibits period doubling, with $M > 0$ for even periods and $M < 0$ for odd periods. Second, the decay of this period-doubling behavior is directly controlled by the frequency of the drive. Third, the lifetime of the time-

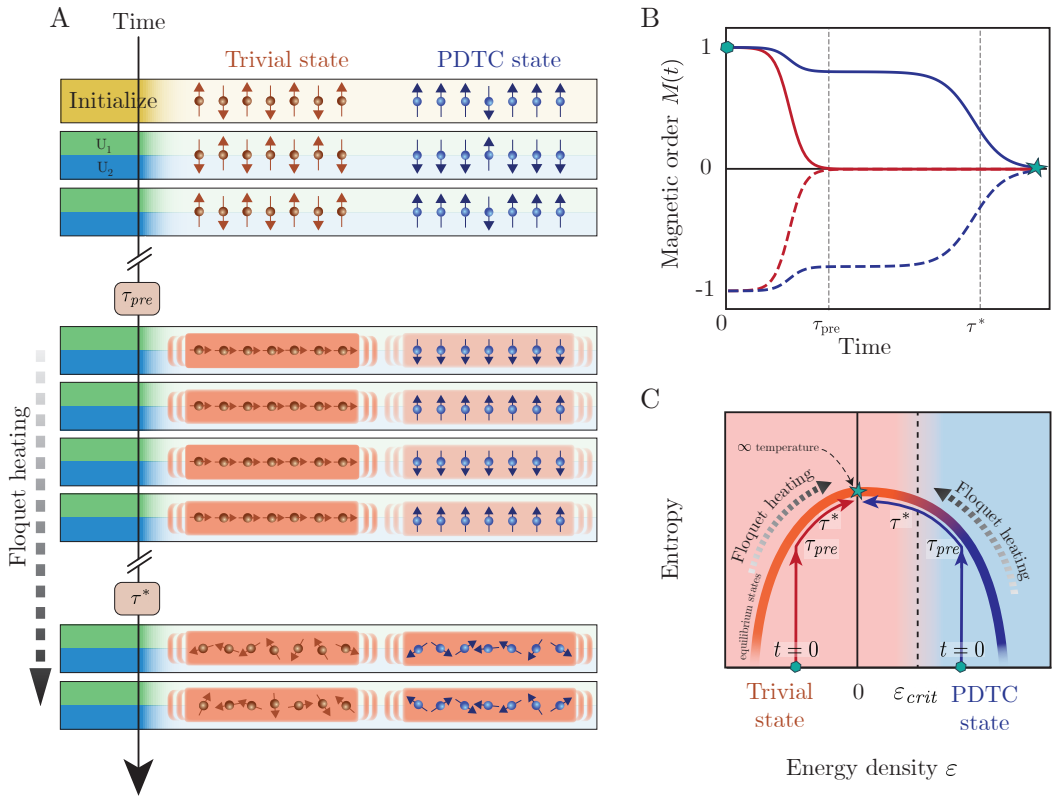

Figure 4.6: **The symmetry-breaking and the symmetry-preserving regimes. A.** For intermediate times between τ_{pre} and τ^*, the system approaches an equilibrium state of the prethermal Hamiltonian H_{eff}. After τ_{pre}, the magnetization in the trivial Floquet phase remains constant. Meanwhile, in the PDTC phase, the magnetization oscillates each period leading to a robust sub-harmonic response. For both phases, at times $t \gg \tau^*$, Floquet heating eventually brings the many-body system to a featureless infinite temperature ensemble. **B.** Schematic of the stroboscopic magnetization dynamics in the trivial (red) and PDTC (blue) phase (full/dashed curves represent even/odd driving periods). In the trivial phase, any transient time-crystalline-order decays by the prethermal equilibration time τ_{pre}, while in the PDTC phase, the order remains robust until τ^*, the frequency-controlled heating timescale. **C.** Starting from a product state with zero entropy, the dynamics under H_{eff} bring the system to an equilibrium state at time τ_{pre}. The PDTC behavior is robust if the initial state thermalizes to a prethermal equilibrium state, which spontaneously breaks an emergent symmetry of H_{eff}. In our system, this occurs if the energy density of the initial state is above a critical value $\varepsilon_{\text{crit}}$ (i.e. to the right of the dashed line), wherein the system equilibrates to a ferromagnetic state. Regardless of the initial state, for $t > \tau^*$, Floquet heating eventually brings the system to the maximum entropy state at zero energy density. Figure adapted from [37].

crystalline order mirrors the dynamics of the energy density shown in Fig. 4.4B, demonstrating that Floquet heating ultimately melts the PDTC at late times.

By considering two additional initial states, we explore the stability of the PDTC phase as a function of energy density. Fig. 4.7C depicts both the heating time as well as the lifetime of the time-crystalline order. Near the bottom of the spectrum, where no symmetry-breaking phase exists, the decay of the magnetization is both frequency-independent and significantly faster than the heating timescale. In contrast, near the top of the spectrum, where the symmetry-breaking ferromagnetic phase lies, the two timescales are consistent with one another, demonstrating that the PDTC lifetime is limited by Floquet heating. Our results are consistent with a phase boundary occurring around energy density $\langle H_{\text{eff}} \rangle / (N J_0) \approx 2$, in agreement with independent numerical calculations via quantum Monte Carlo [37].

4.6 Error sources

The dynamics observed in this work are the combination of ideal Hamiltonian evolution as in Eq. 4.1 and other terms of smaller magnitude that we refer to as "error sources". The combined effect of the latter, when measuring the chain magnetization, manifests as decoherence.

The most significant error source is fluctuating AC Stark shifts of the hyperfine qubit frequency. This fluctuation is mostly caused by power instability of the 355 nm laser light at the ions' location. Even though there is a power PI locking scheme in effect for the 355 nm light, the sampling point for the lock is at a more upstream

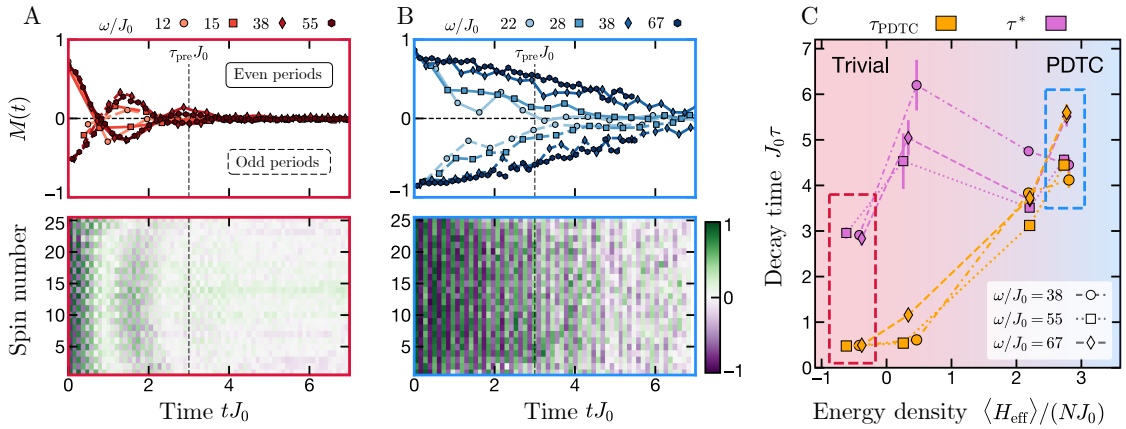

Figure 4.7: **Characterizing the PDTC phase. A, B.** Upper plots: The magnetization dynamics, $M(t)$, for the Néel state (left) and the polarized state (right). For the Néel state, $M(t)$ quickly decays to zero at time $\tau_{\rm pre}$ (dashed vertical line), independent of the drive frequency. For the polarized state, the sub-harmonic response ($2T$-periodicity) persists well-beyond $\tau_{\rm pre}$ and its lifetime is extended upon increasing the drive frequency. The lifetime of the prethermal time-crystalline order, $\tau_{\rm PDTC}$, is extracted from a fit of the magnetization evolution to a decaying exponential. Statistical error bars are of similar size as the point markers. Lower plots: The X-magnetization dynamics for each ion in the chain at $\omega/J_0 = 38$. **C.** Heating (τ^*) and magnetization decay ($\tau_{\rm PDTC}$) times for four different initial states at varying energy densities. For low energy densities, $\tau_{\rm PDTC}$ (orange) are short, independent of frequency, and significantly shorter than τ^* (magenta), highlighting the trivial Floquet phase. For high energies, $\tau_{\rm PDTC}$ is similar to τ^*, highlighting the long-lived, frequency-controlled nature of the PDTC behavior. The location of the observed crossover in energy density is in agreement with an independent quantum Monte Carlo calculation (red and blue shaded regions) [37]. Error bars for the decay time correspond to fit errors, while error bars for the energy density correspond to statistical errors. Figure adapted from [37].

location than the ions. As the beams propagate downstream from that point, active elements, acoustic noise, and air turbulence introduce extra power noise. At the ions location, the light's red and blue beatnotes ideally produce exactly opposite AC Stark shifts of the qubit levels and cancel each other. In practice, however, they are not always perfectly balanced. In this case, common-mode power fluctuations will make the sum of their Stark shifts fluctuate. This manifests as an effective fluctuating magnetic field term $B^{(AC)}(t)\sum_i \sigma_i^z$, common for all spins i and is present

in every stage of the experimental sequence. A fortunate side effect of the π-rotations of the drive is that they echo out part of this noise. However, the spectral portion of $B^{(AC)}(t)$ that is faster than $\omega/2$ is not echoed out and differs between different repetitions, manifesting as decoherence in the final averaged signal. In numerics presented in the next section, we model this noise based on experimental evidence and reasonable simplifications, and present numerical simulations that include it.

Imperfect qubit state readout also impacts the final fidelity of the simulation. During the finite readout window of 400μs, there is a small probability that a $|\downarrow\rangle$ state will be off-resonantly pumped, and read out as a $|\uparrow\rangle$ state, and vice versa. For the experiments presented in this work, the average readout error was 2.3% for each ion.

Another error source comes from a term combining the spin and the motional part of the qubit wavefunction that acts in parallel with the effective Ising interaction. This term represents entanglement between these two parts; when we measure the qubit spin, we effectively trace out the entangled motional state, resulting in a probabilistic mixed state. The probability for such an erroneous spin flip to occur is proportional to

$$\sum_{m=1}^{N} \left(\frac{\eta_m b_{im} \Omega}{\delta_m} \right)^2 \tag{4.4}$$

Therefore, by increasing this detuning, we minimize the undesired spin-motion entanglement, but we are also decreasing the strength of our spin-spin interaction term (see Supplement, Simulating the transverse field Ising Hamiltonian). We set the balance between these effects by keeping the sum in (4.4) less than 0.1 for two spins,

which for the 25 spins results in approximately 0.7% flip probability per spin. See 3.1 for further details. This effect is somewhat amplified by the finite duration of the Hamiltonian quenches in the second term of the Floquet drive, whose spectral decomposition has nonzero components in the motional frequencies. We considerably mitigate this effect by applying the Tukey window shaping to the relevant pulses, which reduces these undesired spectral terms.

Chapter 5: Ion heating from electric field noise

5.1 Overview

The ions are charged particles and as such, any electric field that's not part of the quantum simulation sequence is noise that exerts a force on them and classically heats them. There is a number of electric field noise sources [41], such as fluctuating patch potentials on the trap electrodes. It is commonly known that the heating rate from this noise takes a simple form for harmonic confining potentials (see next section). What happens when the potential is not harmonic to begin with, or when the ion has heated out of the harmonic regime of that potential?

So far, electric field noise has not been something to worry much about in quantum simulators. Ions are Doppler-cooled at the beginning and at the end of an experimental sequence, "resetting" their motional energy very close to zero. And during the sequence itself, other sources of decoherence are much more significant. It is, however, an interesting topic to explore. And perhaps microscopic ion traps of the future might feature slots with shallow confining potentials where ions are stored without Doppler-cooling to avoid scattered light. Or measured ion heating rates might be used to estimate the noise spectrum they see.

In this chapter, I will present proof-of-principle reasoning and numerical cal-

culations. The key points are a) determining the "filter" function selecting the noise frequency components causing heating, b) gaining intuition about how ions heat differently depending on how the confining potential differs from a harmonic one, and c) a de facto demonstration of a method for generating a stochastic signal with a desired spectrum, used in the numerics.

This is a work in progress. I do not aim to provide a fully realistic simulation of the ion heating, and therefore I do not account for more than one spatial dimension, more than one ion, or the effect of the electric field noise *gradient* in parametrically heating the ion, to name a few.

5.2 Heating in a harmonic potential

A generic confining potential can be approximated to be harmonic at the vicinity of a local minimum. This property stems from keeping the quadratic term in its Taylor expansion and ignoring all higher orders. Ions trapped in such potentials with secular frequency ω, when subject to stochastic electric field noise, will in principle heat classically. The relation between this heating rate and the spectrum of the electric field noise takes a simple form [42, 41, 43]:

$$\dot{\bar{n}} = \frac{q^2}{4m\hbar\omega} S_{\mathcal{E}}(\omega) \tag{5.1}$$

where $\dot{\bar{n}}$ is the heating rate of a single ion in motional quanta/s, q and m are the particle charge and mass respectively, and $S_{\mathcal{E}}(\omega)$ is the PSD of the noise in electric field $\mathcal{E}$ for the trap frequency ω. We can re-write (5.1) in terms of the energy change

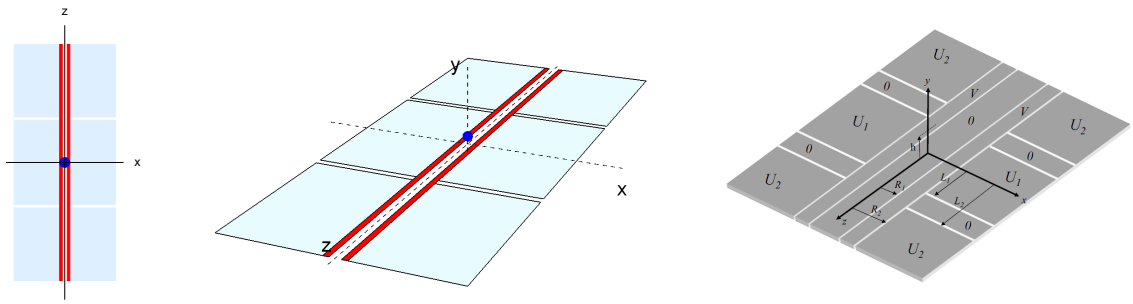

Figure 5.1: **Schematic of a simple chip trap. Left & center:** Top and 3D view of the trap. The red electrodes carry RF potentials, while the blue ones are at fixed voltages. The ion (blue dot) floats slightly above the trap, at a height h of the order of a few tens of μm. **Right:** Symbol notation. From [].

rate $\dot{E} = \dot{\bar{n}}\hbar\omega$ and $S_F(\omega)$, the PSD of the random force F exerted on the particle by the field:

$$P_0 = \dot{E} = \frac{S_F(\omega)}{4m} \tag{5.2}$$

where I used $S_{\mathcal{E}} = q^2 S_F$.

In plain words, what (5.2) tells us is that the only frequency component of the noise that causes heating is the one that is resonant with the trap frequency. Naturally, one asks: is the same true for confining potentials that are anharmonic and the particle's oscillation frequency depends on its total energy E, unlike the harmonic case? To start working on this question, let's first get an idea of what a potential of interest might look like.

5.3 What potential do the ions live in?

Macroscopic ion traps such as rod or blade traps typically create a confining potential that is deeper then 1 electron Volt [44, 45]. Depending on the complexity of their structure, a harmonic potential may be a sufficient approximation up to a

fraction of this depth, or dedicated electrodes may be used to shape the bottom of the potential to an anharmonic form. A quartic potential for example balances out the Coulomb repulsion between ions in the chain in a way that results in uniformly spaced chains [45].

However, the future of quantum information-oriented ion traps seems to be lying with smaller, shallower structures. Their potentials are a fraction of an electron Volt deep (tens of meV), meaning that the harmonic approximation at the bottom breaks down sooner, as ions heat. As such, heating at the anharmonic regime is more relevant for chip traps than macroscopic ones. A basic chip trap with two RF rail electrodes providing transverse confinement and static voltage electrodes providing axial confinement is shown in Fig. 5.1. Using the Biot-Savart-like law for the electric field [46], we can calculate the electric fields and therefore the electric potential along any axis above the trap. For the example trap parameters shown in Table 5.1, these potentials are shown in Fig. 5.2. For the X and Y axes, I calculate the oscillating electric field at the RF frequency Ω_{RF} and then derive the time-averaged field and the resulting static pseudopotential following [23].

Table 5.1: Model chip trap parameters

$L_1 = 250$	$L_2 = 270$	$L_3 = 700$
$R_1 = 15$	$R_2 = 30$	$R_3 = 300$
$U_0 = 5$	$U_1 = -15$	$U_2 = -18$
$V_0 = 40$	$\Omega_{\mathrm{RF}} = 2\pi \cdot 25$	

[1] Lengths are given in μm, voltages in V, frequency in MHz.

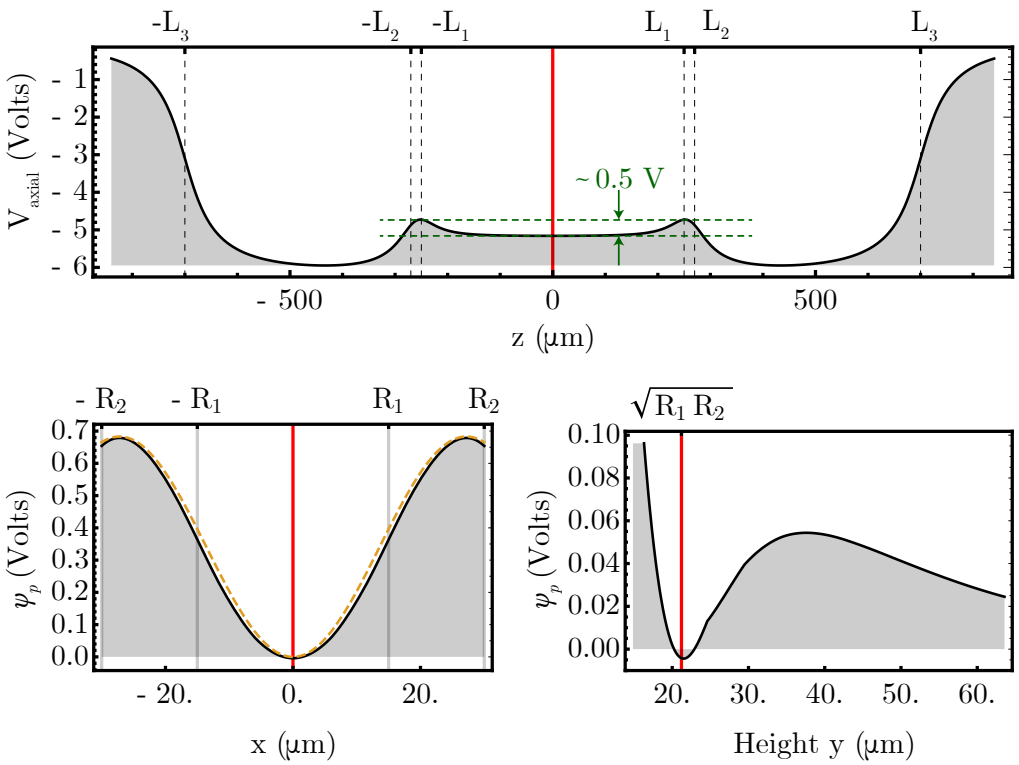

Figure 5.2: **Potentials along all 3 axes of the chip trap from Fig. 5.1.** The red line indicates the ion equilibrium position at each axis. **Top:** Potential along the ion chain axis (Z). The ion chain is centered in the central valley. **Bottom left:** Along the transverse axis X (parallel to the trap surface), the static pseudopotential is symmetric, $\psi_p(x, y, z) = \psi_p(-x, y, z)$. An inverted cosine (dashed orange) is a good approximation, making it equivalent to the simple pendulum problem. **Bottom right:** Along the transverse axis vertical to the trap chip, the static pseudopotential is not symmetric. The ion equilibrium position is at height $h = \sqrt{R_1 R_2}$. The potential depth is about 50 meV for the parameters of Table. 5.1.

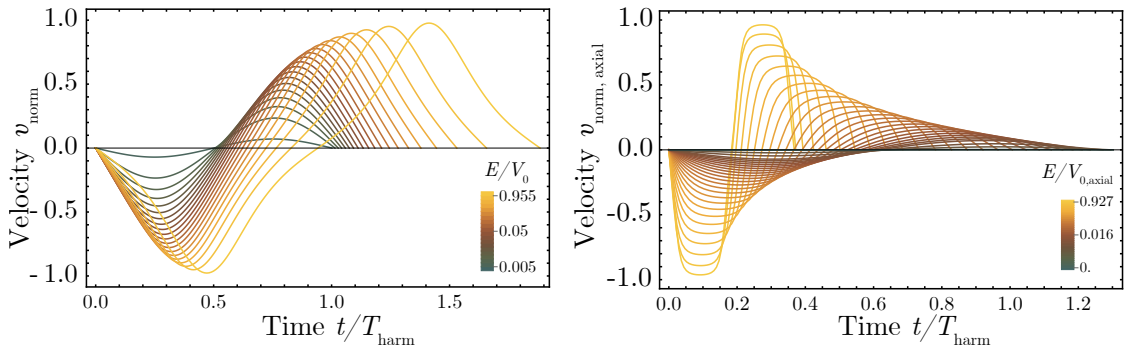

Figure 5.3: **Velocity for periodic motion in the transverse (X) and axial (Z) potentials for a single period. Left:** The velocity $v_{\text{norm}}(t)$ for one oscillation along the transverse X axis, as the ions heats and its energy is increased (dark to lighter curves). At the beginning, when its energy E is low compared to the potential depth V_0, the harmonic approximation of the potential is accurate and its motion is harmonic: $v_{\text{norm}}(t)$ is a sinusoidal and lasts exactly T_{harm}. As its energy is increased though, the motion becomes less and less sinusoidal, and lasts longer. The velocity is normalized to its value for the maximum energy V_0, $\sqrt{2V_0 e/m}$. The oscillations last longer with higher energies, because the potential widens compared to the harmonic approximation and the restoring force experienced is on average less than in the harmonic case. **Right:** Same plot for the axial potential. Contrary to the case on the left, now as the ion heats, its oscillations become faster. This is because for the parameters chosen, the potential at higher energies becomes narrower than its harmonic approximation, and the restoring force is stronger. Note that the energy color scale is nonlinear.

5.4 Oscillation period and susceptibility to noise

It is clear that the harmonic approximation is only valid for a small fraction of the potential depth. Particles in non-quadratic potentials perform periodic, but not harmonic motion. To provide additional intuition about the effect of anharmonicity to particle motion, I calculate classical ion trajectories at the absence of noise and show the velocity during a single period in Fig. 5.3. In general, depending on whether the potential gets narrower or wider than its harmonic approximation at high energies, the oscillation period will become shorter or longer respectively, as the ion heats.

Looking at the transverse velocity plot (Fig. 5.3, left), we can start visualizing that as the ion heats and its oscillation slows down, noise frequencies that were resonant earlier now are becoming too fast and less relevant, similar to an oscillator driven off-resonance by frequencies faster than its natural one. Unfortunately, realistic noise spectra are noisier at lower frequencies, meaning that the heating rate likely increases as the ion heats.

On the contrary, at the axial case (Fig. 5.3, right), as the ion heats, its motion becomes resonant with faster frequencies. This is good news, since noise typically decreases with frequency in typical spectra.

In the following section, I will show how we can use the velocity calculations of Fig. 5.3 to calculate an effective filter on the electric field noise, determining which frequency components of it will resonate and amplify the ion motion.

5.5 Anharmonic potentials and noise frequencies that heat

In this section I will formalize the intuition that the noise frequencies causing heating are mostly the ones resonant with ion motion, for a generic potential. I am following the framework set at [47]. The most important assumption is that the noise is sufficiently small, so that over a large number of classical oscillations the ion trajectory during a period remains what it would be at the absence of noise. In other words, noise slowly adds energy, but the trajectory is still described by noiseless evolution at a given energy (which increases with time).

At time t, the infinitesimal change in the particle's energy E is $\mathrm{d}E = f(t)\mathrm{d}x$,

where $f(t)$ is the stochastic noise force and $\mathrm{d}x$ the infinitesimal displacement; the effect of the restoring force from the potential is omitted since it is conservative and it averages out every period. We can re-write $\mathrm{d}E$ as

$$\mathrm{d}E = f(t)v(t)\mathrm{d}t \tag{5.3}$$

where $v(t) = \mathrm{d}x/\mathrm{d}t$ and get for a non-infinitesimal time interval $\Delta t = [t, t + \Delta t]$:

$$\Delta E = \int_t^{t+\Delta t} f(t')v(t')\mathrm{d}t' \tag{5.4}$$

We will choose Δt to be much larger than one period, but small enough that the energy change ΔE is small compared to E. According to this "small noise assumption", we can replace the velocity $v(t)$ with the noiseless evolution velocity $v_0^{(E)}(t)$. The latter is only sinusoidal for a harmonic potential, and may even by asymmetric during one period, depending on the potential profile. It can be computed numerically for an arbirtrary confining potential. Within Δt though, it is necessarily periodic, and as such we can use its Fourier series expansion:

$$v_0^{(E)}(t) = \sum_{k\neq 0}^{\infty} a_k(E)e^{-ik\omega(E)t} \tag{5.5}$$

where both the coefficients $a_k(E)$ and the basis functions of the expansion depend on the energy E at the start of the interval Δt. It is essentially $v_0^{(E)}(t)$ that is shown in the plots of Fig. 5.3 for different values of energy. The coefficients $a_k(E)$ depend only on the confining potential $V(x)$, they are a function of energy E, and are given

by

$$a_k(E) = \frac{1}{T(E)} \int_0^T v_0^{(E)}(t) \cdot e^{ik\omega(E)t} \mathrm{d}t \tag{5.6}$$

Substituting (5.5) back into (5.4), we get

$$\Delta E(E) = \sum_{k \neq 0}^{\infty} a_k(E) \cdot F_{t,\Delta t}\left[k\omega(E)\right] \tag{5.7}$$

where $F_{t,\Delta t}\left[k\omega(E)\right]$ is the truncated Fourier transform of the stochastic force $f(t)$ between times t and $t + \Delta t$:

$$F_{t,\Delta t}\left[k\omega(E)\right] = \int_t^{t+\Delta t} f(t') e^{-ik\omega(E)t'}$$

As a final tweak, we can use the fact that $a_k(E)$ and $F_{t,\Delta t}$ are essentially Fourier transforms of real quantities and thus even-symmetric at the frequency axis, $a_{-k}(E) = a_k^*(E)$ and similarly, $F_{t,\Delta t}\left[-k\omega(E)\right] = F_{t,\Delta t}^*\left[k\omega(E)\right]$:

$$\Delta E(E) = 2 \sum_{k>0}^{\infty} \mathrm{Re}\left\{a_k(E) F_{t,\Delta t}\left[k\omega(E)\right]\right\} \tag{5.8}$$

Let's unpack what (5.8) tells us. For a given energy E, the energy change ΔE is directly related to the noise frequency resonant with the frequency $\omega(E)$ at this energy and its harmonics $k\omega(E)$, weighted by the factors $a_k(E)$. *The profile of $a_k(E)$ with frequency $k\omega(E)$ acts as a filter acting on the noise spectrum, determining the efficiency with which the noise frequency components lead to heating.* In Fig. 5.4 I illustrate the profile of $|a_k(E)|^2$ for various energy levels at the transverse potential

along X (from Fig. 5.2).

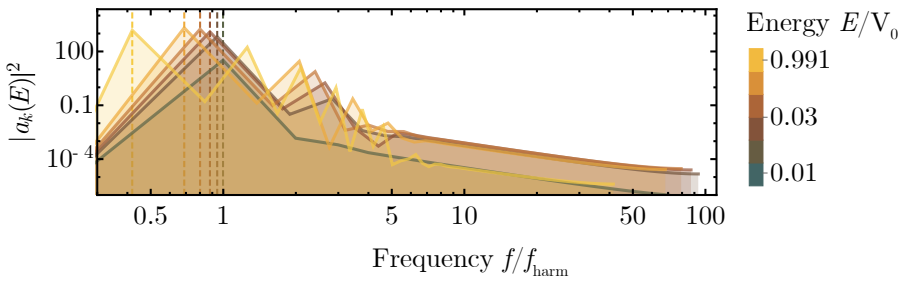

Figure 5.4: $a_k(E)$ **factors for various energies** E **for the transverse** (X) **potential.** The frequency axis is normalized to f_{harm}, the frequency corresponding to the harmonic approximation of the potential. For low energies (darker curves), the harmonic approximation is a good description, manifesting as a peak at $f/f_{\text{harm}} = 1$: noise frequencies resonant with f_{harm} are the main contributors to heating. As the oscillation energy E increases though, the particle motion becomes more complicated (and slow), allowing lower noise frequencies to cause heating.

5.6 Generating signals with a desired spectrum

Now that we have an idea of what to expect for different confining potentials and noise spectra, I will go over the last step needed before running a numerical simulation: generating the stochastic electric field force time series with a spetrum imitating realistic lab conditions. Let's say we want to generate a signal $s(t)$ with a desired power spectral density (PSD) $S(f)$. For infinitely long signals, $S(f)$ starts from infinitesimal frequencies and goes up to infinite frequency. Realistic signals though have a finite duration T_{inst}. This sets the slowest frequency the Fourier decomposition can contain, $f_{\text{min}} \equiv 1/T_{\text{inst}}$. The PSD calculated for this finite signal is the *truncated* PSD $S_{T_{\text{inst}}}(f)$ (see (5.10)).

Additionally, for ease of numerical methods, I will be working with discrete signals, assumed to represent realistic continuous signals sampled at interval dt. The

sampling interval must be chosen to be short enough so that the simulation does not "miss" critical events, but long enough that we do not require a prohibitive amount of computational resources. I typically set dt so that an oscillation of the ion inside the confining potential lasts a few tens of dt. At the same time, dt sets the highest frequency of the Fourier decomposition to $f_{\max} \equiv 1/dt$. Remembering that real (not complex) signals have symmetric Fourier transforms around $f_{\max}/2$, the highest frequency we can have meaningful information for is $f_{\max}/2$. This is why PSDs shown in this chapter will be even symmetric; I include the symmetric part for clarity.

After clarifying the relation of duration to frequency, let's establish a working convention for generating a signal "with a desired spectrum". The PSD is a somewhat abstract concept, since it involves averaging an infinite number of instances (realizations) of the actual signal to specify it. No instance by itself will yield the exact same PSD as another one. For the purposes of this chapter, when setting a goal to generate n instances of a signal in time with duration T_{inst} each, I will be generating a signal nT_{inst} long with a PSD *exactly equal* to the desired one. Then, I will be partitioning it into n pieces, with PSDs that fluctuate around the ideal/desired one. This subtlety is illustrated in Fig. 5.5.

Next, in order to explain the generation of the time signal from a PSD, let's remember the Wiener-Khinchin theorem. It allows defining a (truncated) PSD

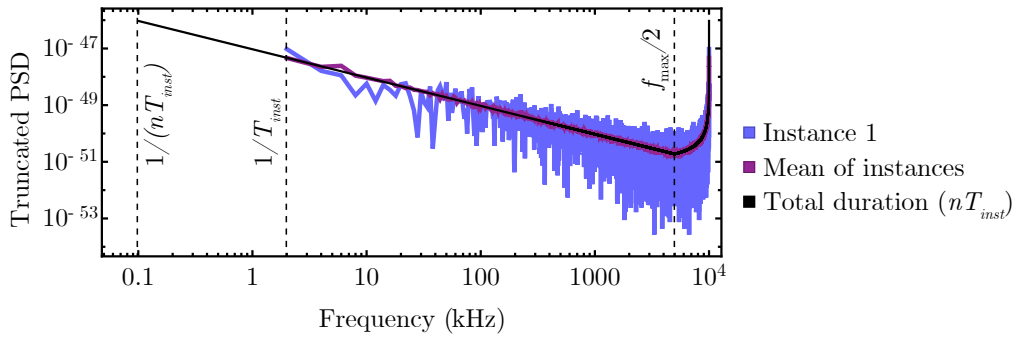

Figure 5.5: The total, nT_{inst}-long time signal generated has exactly the PSD requested (black), here with a $1/f$ dependence. An instance chosen at random, eg. the first (blue) has a PSD which is still $1/f$ but noisy. The averaged PSD of all instances converges to the requested $1/f$ dependence (purple). The PSD for the total signal is calculated from a longer sequence compared to the PSD for a single instance, thus it extends to lower frequencies.

$S_{T_{\text{inst}}}(\omega)$ through the (truncated) Fourier transform $F_{T_{\text{inst}}}(\omega)$ of a signal $s(t)$:

$$F_{T_{\text{inst}}}(\omega) = \int_0^{T_{\text{inst}}} e^{-i\omega t} s(t) \mathrm{d}t \tag{5.9}$$

$$S_{T_{\text{inst}}}(\omega) = \frac{1}{T_{\text{inst}}} \langle |F_{T_{\text{inst}}}(\omega)|^2 \rangle \tag{5.10}$$

According to our working convention from earlier, the desired form $S_{\text{des}}(\omega)$ for the PSD will be made to exactly correspond to the spectrum of the nT_{inst}-long signal:

$$S_{\text{des}}(\omega) = \frac{1}{nT_{\text{inst}}} |F_{nT_{\text{inst}}}(\omega)|^2 \tag{5.11}$$

and therefore we can write the Fourier transform of the total signal as

$$F_{nT_{\text{inst}}}(\omega) = \sqrt{nT_{\text{inst}} S_{\text{des}}(\omega)} e^{i\phi(\omega)}. \tag{5.12}$$

The phase $\phi(\omega)$ is the randomness-adding element of this algorithm. Working with

actual discrete spectra and signals in the simulation, $\phi(\omega)$ becomes ϕ_k for each discrete frequency ω_k and is sampled from a uniform distribution between 0 and 2π. Now that we have calculated the Fourier transform $F_{nT_{\text{inst}}}(\omega)$ of the signal $s(t)$, all that is left to do is apply the inverse Fourier transform F^{-1} to it, to transition to the time domain:

$$s(t) = \int_{\omega_{\min}}^{\omega_{\max}} F_{nT_{\text{inst}}}(\omega) e^{i\omega t} \mathrm{d}\omega \tag{5.13}$$

5.7 Numerical results for a harmonic and an anharmonic potential

I start with the case of a truncated harmonic potential $U_{\text{harm}}(x)$ and simulate the ion behavior under a noisy electric field. Using a quadratic potential allows to benchmark the numerics against analytical expectations for the heating rate. I then proceed to an anharmonic potential $U(x)$ of width w and depth U_0, equal to the truncated sinusoid that has U_{harm} as its approximation at the bottom (Fig. 5.6).

$$U_{\text{harm}}(x) = \frac{1}{2} m (2\pi f_0)^2 x^2 \tag{5.14}$$

$$U(x) = \frac{1}{2} U_0 \left(1 - \cos \frac{\pi x}{w}\right) \tag{5.15}$$

where f_0 is the secular frequency of the harmonic regime, and $U_0 = 8 f_0^2 w^2 m$ to make sure that $U(x) = U_{\text{harm}}(x) + O(x^4)$ around $x = 0$. $f_0 = f_{\text{harm}}$ is the secular frequency of the harmonic approximation.

The code used numerically solves the differential equation of motion

$$\frac{\mathrm{d}^2 x}{\mathrm{d}t^2} = -\frac{\mathrm{d}U_{\text{pot}}(x)}{\mathrm{d}x} + f(t)$$

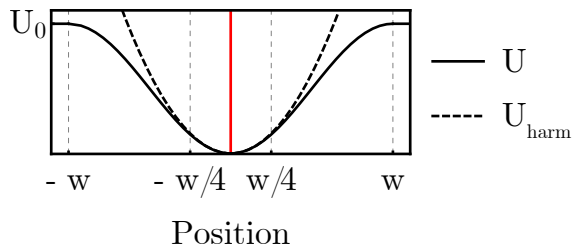

Figure 5.6: **Confining potentials used for the simulation** The ion equilibrium position is at the red line.

where $x(t)$ is the ion's position, $U_{\text{pot}}(x) = U(x)$ or $U_{\text{harm}}(x)$ is the confining potential, and $f(t)$ is the stochastic force from the noisy electric field. Using the method from 5.6, $f(t)$ can be made to have an arbitrary spectrum. In all subsequent simulation results, the PSD of the noise falls as its inverse frequency, a common occurence in the lab [48, 49]:

$$S_F(\omega) = S_0 \frac{2\pi f_0}{\omega}$$

In principle though it may be an analytical function or an experimentally measured spectrum.

I discretize the signals with a time step dt set to 1/10 of the ion period at the harmonic regime. Also, because of limited computational resources, the parameters used for this simulation do not correspond to realistic ion traps. Specifically, the chosen heating rate (determining S_0) is chosen as 13000 quanta/s, much higher than typical numbers of a few tens to hundreds of quanta/s. Additionally, the width $w = 2$ μm of the confining potential is much smaller than typical widths. These choices, however, allow for significant heating within a duration that can be simulated in a few minutes with our resources. Using a more advanced system, one can afford to get much more realistic and computationally demanding. The total generated

signal for each potential case was 2 s, with $dt = 1$ μs, split into $n = 50$ instances. The harmonic frequency was $f_0 = 0.1$ MHz. Fig. 5.7 shows the simulation results for a particular noise instance and an ion that began from resting at the bottom U and U_{harm}.

From the position and velocity of the ions I calculate their energy evolution for each instance at Fig. 5.8. Despite the small number of instances used (50), their ensemble averages already tell different stories for each potential, confirming our expectations from previous sections: it seems that the heating rate for the anharmonic case becomes higher than for the harmonic case after some time, corresponding to the average time needed for the ion to heat out of the harmonic approximation regime. Notice that as the energy of the ensemble increases, so does its variance.

5.8 The Fokker-Planck equation

The distribution of the ion energy E as a function of time, $\eta(E, t)$, can be modeled with the Fokker-Planck equation formalism [47]. Essentially, this formalism models the evolution of a probability distribution (of the energy in our case) under the action of a stochastic process:

$$\frac{\partial \eta}{\partial t} = -\frac{\partial (g_1 \eta)}{\partial E} + \frac{1}{2}\frac{\partial^2 (g_2 \eta)}{\partial E^2}$$

where $g_1 = g_1(E)$ is the drift coefficient, characterizing the ensemble heating rate, and $g_2 = g_2(E)$ is the diffusion coefficient, characterizing the spread of the energy distribution as the ensemble heats. I will leave a theoretical treatment for the future,

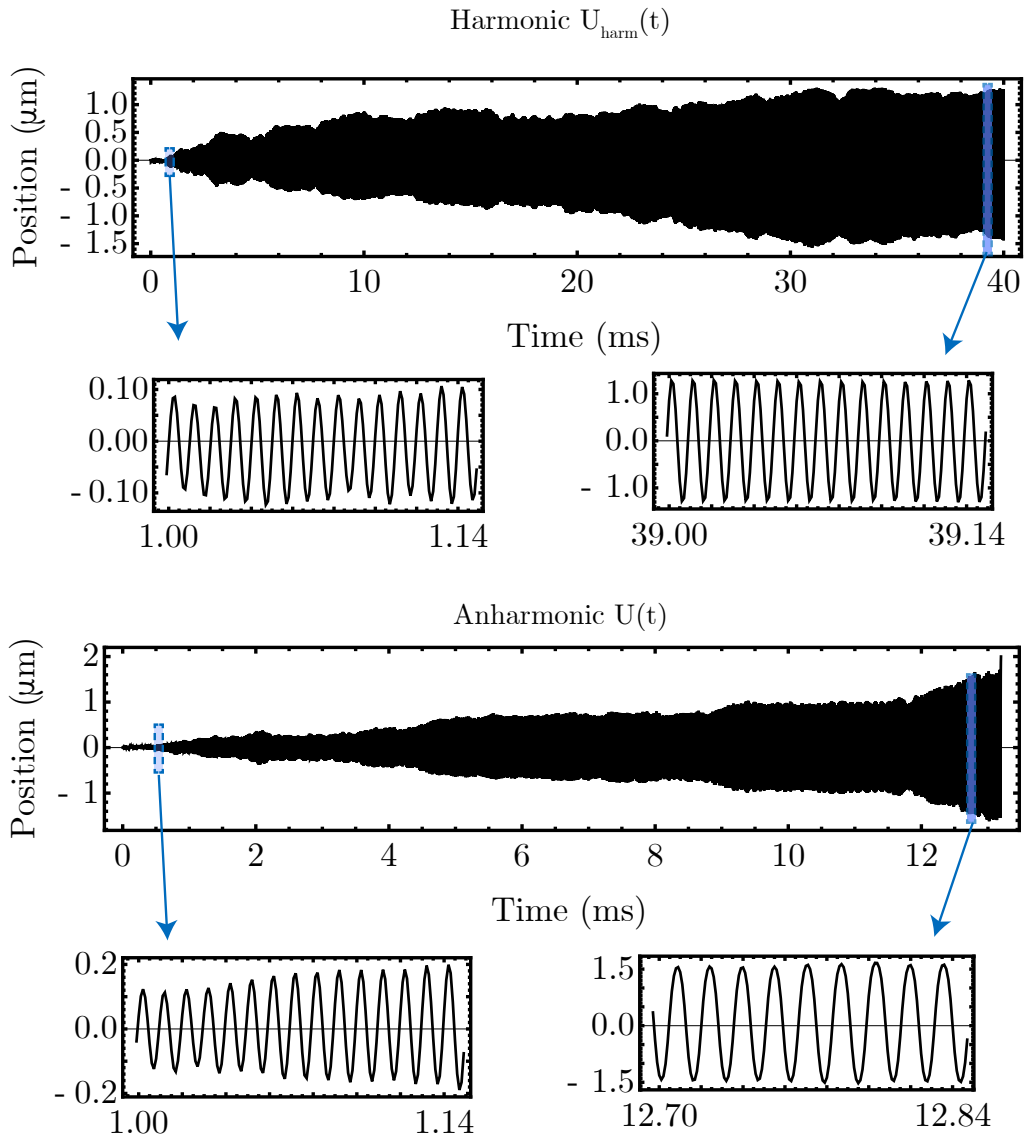

Figure 5.7: **Position $x(t)$ for the harmonic and the anharmonic potential.** Shown is an instance chosen at random for each. Not many conclusions can be deduced from single instances; there are howver two comments to be made. First, notice that the ion oscillation for the anharmonic case (bottom) has visibly slowed down at 12 ms, since at higher energies this potential hosts slower oscillations. Second, the time trace for the anharmonic case is shorter than for the harmonic, because around 13 ms the particle escaped the potential: its energy reached $\approx U_0$ and its motion amplitude $\approx w$.

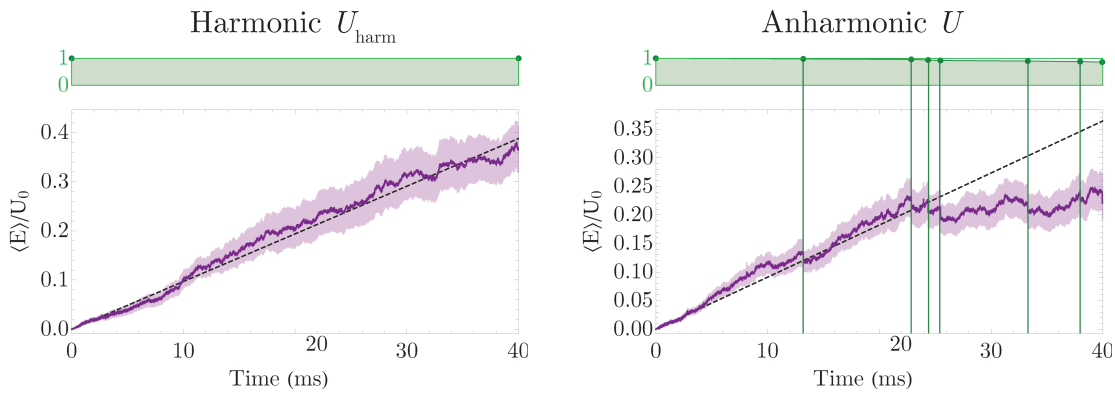

Figure 5.8: **Energy evolutions for the harmonic and anharmonic potentials** The ensemble average for the 50 instances at each time step is shown in purple. The band denotes the standard error of the mean, while the black dotted line shows linear heating from the fixed harmonic heating rate P_0 from (5.2). The harmonic case seems to follow that prediction. The anharmonic case also seems to follow it at the beginning and until ~ 5 ms. After that, we have recurring events when the ensemble average heats faster than P_0, until an ion in some instance heats so much it escapes the trap (vertical green lines). Right after, the ensemble average of the remaining instances soon picks up again. The green plots trace those "escape events" by showing the percentage of ions staying trapped at time t.

but it is worth noting that both the drift and the diffusion coefficients are related to the numerically calculated energy distribution, shown in Fig. 5.9. Additionally, Fig. 5.10 shows the *variance* of the ΔE added during an interval Δt, as a function of the ion energy E over that interval Δt. In both the harmonic and the anharmonic case, this relationship is linear. This variance $\sigma^2(\Delta E)$ is directly related to the diffusion coefficient $g_2(E)$, which can therefore also be numerically estimated.

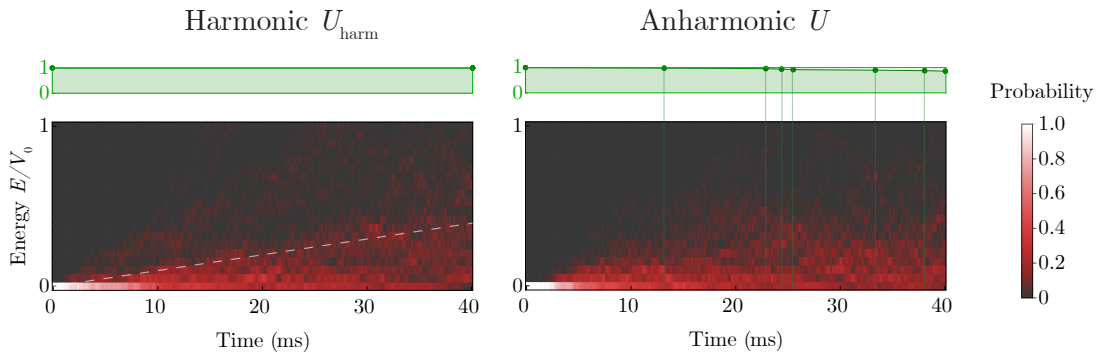

Figure 5.9: **The evolution of the energy distribution for the harmonic and the anharmonic potential.** Each vertical slice is the distribution of the energies for that time step. At $t = 0$, all ions start from zero energy, and depending on the stochastic $f(t)$ for each instance, they heat with different rates. The dashed white line in the harmonic case is the ensemble averaged $\langle E \rangle(t)$ and can be used to numerically find the drift coefficient g_1.

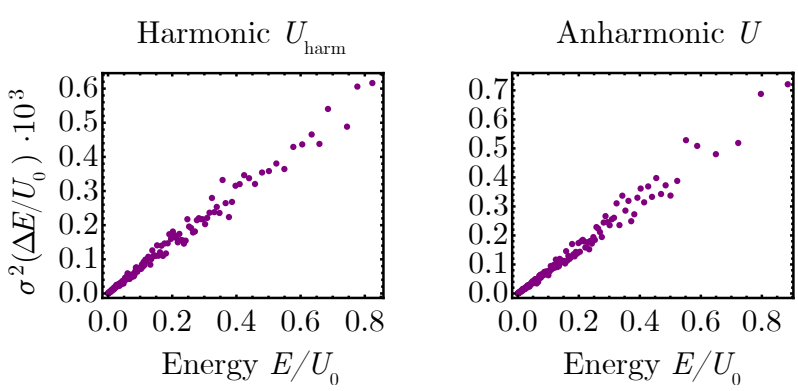

Figure 5.10: **The variance σ^2 of the energy changes ΔE over an interval Δt is linear to the energy E at the beginning of the interval, for both potentials.** To calculate the points shown, I bin $\{E, \Delta E\}$ pairs according to E and then calculate the variance of ΔE within each bin. All energy quantities are normalized to the potential depth U_0.

Appendix A: Stark shifts calculation

The final formulas produced were checked against a Ramsey measurement of the differential Stark shift from the red beatnote and produced a result 5% off the measured value.

A.1 Relevant Stark shifts

The Stark shifts relevant to our operation are 2nd and 4th order (Fig. 2.3). The 2nd order ones stem from the existence of any single Raman beam, which couples the qubit states, $|\uparrow\rangle_z$ and $|\downarrow\rangle_z$, to a subset of the P manifold, and thus shifts each of $|\uparrow\rangle_z$ and $|\downarrow\rangle_z$ with a 2nd order shift. The 4th order ones stem from the coupling of a qubit state to another S state via an intermediate, adiabatically eliminated P state, with the Λ scheme of 2 optical combs (beams).

To calculate Stark shifts, I have found the following prescription helpful:

1. State which state $|i\rangle$ we aim to calculate the Stark shift of.

2. Note the polarization of the light, σ_-, π, and σ_+.

3. List the states $|i'\rangle$ that this light can couple to (Rabi flop), both for single-photon and two-photon Rabi flopping. To do that, use the light's polarization

and selection rules.

4. For each of these states $|i'\rangle$, calculate the detuning of the light from the transition $|i\rangle \leftrightarrow |i'\rangle$, including its sign.

5. Using these detunings and for each of the states $|i'\rangle$, calculate the Stark shift from the base formula (2.9); add them together. Single-photon Rabi flopping will be associated with a 2nd order Stark shift; two-photon (Raman) Rabi flopping will be associated with a 4th order shift.

6. Repeat steps 2 through 5 for all beams. A beam in this case is any light with a unique {intensity, polarization, frequency} triplet. Add the Stark shifts together.

A.2 RF combs

Pulsed lasers like our 355 nm Raman laser create optical combs with a large number of frequencies spaced by the repetition rate $\omega_{\text{rep}} = 2\pi f_{\text{rep}}$. As far as calculation of Rabi frequencies and Stark shifts goes, each of these teeth will behave like a separate CW beam. Therefore, I will review some frequency comb "accounting" before proceeding.

Assume we have two beams, each an optical comb, OC1 and OC2, with a relative detuning of $\delta\omega_{\text{comb}} = 2\pi\delta f_{\text{comb}}$. Their optical comb teeth are given by

$$f_l^{(\text{OC1})} = f_{\text{opt. carrier}} + l f_{\text{rep}} \tag{A.1}$$

$$f_l^{(\text{OC2})} = \delta f_{\text{comb}} + f_{\text{opt. carrier}} + l f_{\text{rep}} \tag{A.2}$$

where $l \in \mathbb{Z}$, $l \in [-l_0, l_0]$, and l_0 is related to the pulsed laser generating the combs and its bandwidth $f_{BW}^{(\mathrm{OC})}$: $l_0 \approx 4 f_{BW}^{(\mathrm{OC})}/f_{\mathrm{rep}} = \frac{4}{\tau f_{\mathrm{rep}}}$, which in QSim's case is $l_0 \approx 1200$[1].

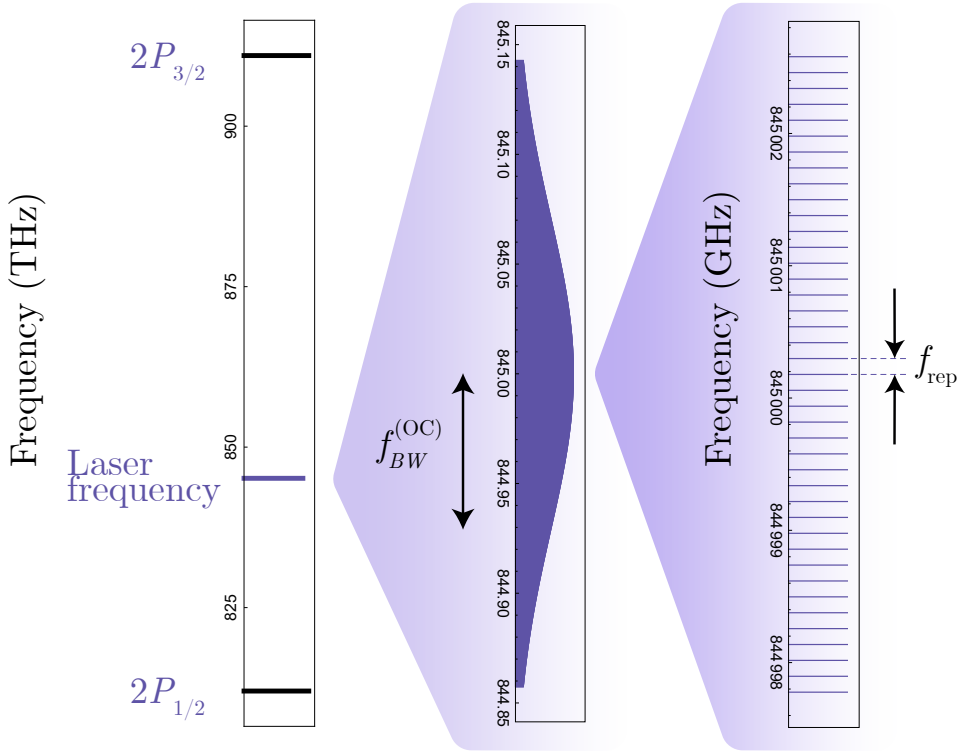

Figure A.1: Optical comb of our 355 nm pulsed laser. The laser carrier frequency $c/355$ nm ≈ 845 THz lies between states ${}^2P_{1/2}$ and ${}^2P_{3/2}$ (left). Zooming in (center), we see the electric field envelope of bandwidth $f_{BW}^{(\mathrm{OC})} \approx 71$ GHz. Zooming in further (right), we see the individual comb teeth, spaced by $f_{\mathrm{rep}} \approx 120$ MHz.

.

In the context of Raman transitions and 4th order Stark shifts, it is useful to look at the beatnote between the two optical combs, which we call an "RF comb" (RFC), since its frequencies lie on the RF frequency spectrum. To find out the frequencies $f_k^{(\mathrm{RFC})}$ of the RF comb teeth, we write down all the combinations of an

[1] I choose to include most of the intensity envelope by setting its width to four times its bandwidth.

optical comb tooth from OC2 subtracted from OC1 and vice versa:

$$\text{RFC from two OCs} = \begin{cases} \text{OC1} - \text{OC2}: & l_1 f_{\text{rep}} - \delta f_{\text{comb}} - l_2 f_{\text{rep}} \\ \\ \text{OC2} - \text{OC1}: & l_2 f_{\text{rep}} + \delta f_{\text{comb}} - l_1 f_{\text{rep}} \end{cases} \tag{A.3}$$

$$\text{RFC from one OC} = l_1 f_{\text{rep}} \tag{A.4}$$

For the case of two optical combs (A.3), setting $l_1 - l_2 \equiv k$, we get for the RF comb teeth:

$$f_k^{(\text{RFC})} = k f_{\text{rep}} \pm \delta f_{\text{comb}}, \qquad k \in [-2l_0, 2l_0] \tag{A.5}$$

where the $\pm$ term is omittted if there is only one optical comb present, as in (A.4). The detunings $\{\delta_k\}$ of these RF comb teeth from a transition f_{tr} responsible for the Stark shift are

$$\begin{cases} \text{One optical comb:} & \delta_k = f_{\text{tr}} - k f_{\text{rep}} \\ \\ \text{Two optical combs:} & \delta_k = f_{\text{tr}} - k f_{\text{rep}} \pm \delta f_{\text{comb}} \end{cases} \tag{A.6}$$

The transition frequency f_{tr} may be the hyperfine splitting $\omega_{\text{HF}}/(2\pi)$, or any other transition within the S manifold, such as $|1,0\rangle \leftrightarrow |1,-1\rangle$. Depending on light polarization, many of these transitions may not be allowed. Depending on the operation, δf_{comb} may be equal to $f_A = (2\pi)^{-1} \mod (\omega_{\text{HF}}/\omega_{\text{rep}})$ for a simple carrier transition, $f_A \pm \omega_{\text{COM}}$ for resonant sideband flopping, or $f_A \pm (\omega_{\text{COM}} + \delta)$ for the M-S scheme. In the latter, δ is the detuning of the laser beatnote from the COM motional mode.

A.3 Single-photon Rabi frequencies

After the detunings, the next ingredient necessary to calculate Stark shifts is the relevant Rabi frequencies. The single-photon Rabi frequency for the transition between states $|i\rangle$ and $|j\rangle$ for light of frequency ω_l is

$$g_{l,ij} = \frac{\vec{\mu}_{ij} \cdot \vec{E}_l}{\hbar} = E_l \frac{\vec{\mu}_{ij} \cdot \hat{\varepsilon}}{\hbar} \tag{A.7}$$

where the electric field corresponding to ω_l is $\vec{E}_l = E_l \hat{\varepsilon}$. Our Raman laser is pulsed, and the existence of pulses in the time domain results to its Fourier decomposition being a frequency comb. The electric field envelope in the pulsed laser is [50]

$$E_{\text{env}}(t) = \sum_{l=-\infty}^{\infty} f(t - lT_{\text{rep}})$$

The Fourier transform of $E_{\text{env}}(t)$ is

$$\tilde{E}_{\text{env}}(\omega) = \int_{-\infty}^{\infty} e^{-i\omega t} E_{\text{env}}(t) \mathrm{d}t \tag{A.8}$$

$$= \int_{-\infty}^{\infty} e^{-i\omega t} \sum_{l=-\infty}^{\infty} f(t - lT_{\text{rep}}) \mathrm{d}t \tag{A.9}$$

$$= \sum_{l=-\infty}^{\infty} \int_{-\infty}^{\infty} e^{-i\omega t} f(t - lT_{\text{rep}}) \mathrm{d}t \tag{A.10}$$

Now imagine that for every term l of the sum, we do the redefinition $t' \equiv t - lT_{\text{rep}} \Rightarrow t = t' + lT_{\text{rep}}$:

$$e^{-i\omega t} f(t - lT_{\text{rep}}) = e^{-i\omega l T_{\text{rep}}} e^{-i\omega t'} f(t')$$

and we have

$$\tilde{E}_{\mathrm{env}}(\omega) = \sum_{l=-\infty}^{\infty} e^{-i\omega l T_{\mathrm{rep}}} \int_{-\infty}^{\infty} e^{-i\omega t'} f(t') \mathrm{d}t' \tag{A.11}$$

$$= \left(\sum_{l=-\infty}^{\infty} e^{-i\omega l T_{\mathrm{rep}}} \right) \cdot \left[\int_{-\infty}^{\infty} e^{-i\omega t'} f(t') \mathrm{d}t' \right] \tag{A.12}$$

$$= \left[\sum_{l=-\infty}^{\infty} 2\pi\delta(\omega T_{\mathrm{rep}} - 2\pi l) \right] \cdot \left[\int_{-\infty}^{\infty} e^{-i\omega t'} f(t') \mathrm{d}t' \right] \tag{A.13}$$

$$= \left[\sum_{l=-\infty}^{\infty} \delta\left(\frac{T_{\mathrm{rep}}}{2\pi/\omega} - l \right) \right] \cdot \left[\int_{-\infty}^{\infty} e^{-i\omega t'} f(t') \mathrm{d}t' \right] \tag{A.14}$$

$$= \sum_{l=-\infty}^{\infty} \int_{-\infty}^{\infty} e^{-il\omega_{\mathrm{rep}} t'} f(t') \mathrm{d}t' \tag{A.15}$$

where I used an exponential sum rule and $\omega_{\mathrm{rep}} = 2\pi/T_{\mathrm{rep}}$. It is clear that the Fourier transform $\tilde{E}_{\mathrm{env}}(\omega)$ is a sum of discrete frequencies (comb teeth). I divide by T_{rep} to get time-averaged value over a pulse cycle for a single comb tooth l:

$$E_l \equiv \frac{1}{T_{\mathrm{rep}}} \tilde{E}_{\mathrm{env}}(l\omega_{\mathrm{rep}}) = \frac{\omega_{\mathrm{rep}}}{2\pi} \int_{-\infty}^{\infty} e^{-il\omega_{\mathrm{rep}} t'} f(t') \mathrm{d}t' \tag{A.16}$$

If $f(t) = \sqrt{\pi/2} E_{\mathrm{env},0} \operatorname{sech}(\pi t/\tau)$ [50], we get

$$E_l = \tilde{E}_{\mathrm{env}}(l\omega_{\mathrm{rep}}) = E_{\mathrm{env},0} \frac{\omega_{\mathrm{rep}}\tau}{\sqrt{8\pi}} \operatorname{sech} \frac{l\omega_{\mathrm{rep}}\tau}{2} \tag{A.17}$$

To use the more easily measurable quantity of time-averaged intensity $\bar{I} \equiv \overline{I_{\text{env}}(t)}$, we note that

$$\bar{I} = \overline{I_{\text{env}}(t)} = \frac{1}{2}c\epsilon_0\overline{E^2_{\text{env}}(t)} = \frac{1}{2}c\epsilon_0\frac{\omega_{\text{rep}}\tau}{2\pi}E^2_{\text{env},0} \tag{A.18}$$

$$\Rightarrow E_{\text{env},0} = \sqrt{\frac{4\pi\bar{I}}{c\epsilon_0\omega_{\text{rep}}\tau}} \tag{A.19}$$

Inserting (A.19) into (A.17), we get

$$E_l = \sqrt{\frac{\bar{I}\omega_{\text{rep}}\tau}{2c\epsilon_0}}\,\text{sech}\,\frac{l\omega_{\text{rep}}\tau}{2}. \tag{A.20}$$

Now we can proceed with the next element to be calculated for (A.7). The matrix elements $\vec{\mu}_{ij}\cdot\hat{\varepsilon}$ are calculated with the aid of Clebsch-Gordan coefficients for the corresponding angular momentum values of $|i\rangle$ and $|j\rangle$; see Jonathan Mizrahi's thesis for details [51].

$$g_{l,ij} = \frac{E_l}{\hbar}C(F, m_F, J, F', m'_F, J', I^{(\text{nucl})}, q)\mu_{ij} \tag{A.21}$$

$$= \frac{E_l}{\hbar}C(\ldots)\sqrt{\frac{3\pi\epsilon_0\hbar c^3\gamma\cdot(2J'+1)}{\omega^3_{ij}}} \tag{A.22}$$

$$= \frac{E_l}{\hbar}\tilde{C}(\ldots)\sqrt{\frac{3\pi\epsilon_0\hbar c^3\gamma\cdot\hbar\gamma/(12\pi c^2)}{\omega^3_{ij}\cdot\hbar\gamma/(12\pi c^2)}}$$

$$= E_l\tilde{C}(\ldots)\gamma\sqrt{\frac{\epsilon_0 c}{4I_{\text{sat},ij}}} \tag{A.23}$$

$$= \frac{\tilde{C}(\ldots)}{2}\cdot\gamma\sqrt{\frac{\bar{I}}{2I_{\text{sat},ij}}}\cdot\sqrt{\omega_{\text{rep}}\tau}\,\text{sech}\,\frac{l\omega_{\text{rep}}\tau}{2} \tag{A.24}$$

$$= g_{ij,0}\frac{\tilde{C}(\ldots)}{2}\sqrt{\omega_{\text{rep}}\tau}\,\text{sech}\,\frac{l\omega_{\text{rep}}\tau}{2} \tag{A.25}$$

Table A.1: Summary table of Stark shift-related parameters

Mode-locked laser pulse duration τ		14.8 ps	
Mode-locked laser pulse rep. rate $\omega_{\text{rep}}/(2\pi)$		120.125 MHz	
$I_{\text{sat},P_{1/2}}$	510.3 W/m^2	$I_{\text{sat},P_{3/2}}$	950.6 W/m^2
Power (each Raman beam)		0.66 W	
Beam waists (ion location, μm)[1]		106.7×8.9 , 114×11.2	
$g_{P_{1/2},0}$	$2\pi \cdot 12.06$ GHz	$g_{P_{3/2},0}$	$2\pi \cdot 12.26$ GHz

[1] As measured on February 3, 2021

where $\tilde{C}(\ldots) \equiv \sqrt{2J'+1}C(\ldots)$, $g_{ij,0} \equiv \gamma_j\sqrt{\bar{I}/(2I_{\text{sat},i,j})}$ and I used the intensity as a function of the electric field, transition saturation intensity [52] $I_{\text{sat},ij} \equiv \pi hc\gamma/(3\lambda_{ij}^3) = \hbar\omega_{i,j}^3\gamma/(12\pi c^2)$. I note that the γ used is the inverse of the spontaneous emission lifetime of about 8 ns [53, 54], for the $P_{1/2}$ state, and about 6 ns for the $P_{3/2}$. Observe that

$$\sum_{l=-\infty}^{\infty} |g_{l,ij}|^2 \approx \int_{l=-\infty}^{\infty} \mathrm{d}l\, |g_{l,ij}|^2 = \tilde{C}^2(\ldots)g_0^2.$$

A.4 Two-photon Rabi frequencies

The discussion so far was relevant to optical comb teeth and single-photon transitions of optical frequencies. For a Raman transition between S levels, say $|0,0\rangle \leftrightarrow |1,0\rangle$, the transition frequency ω_{tr} is in the RF range. It can be that 2 optical comb teeth from the same beam (aka comb) interfere and drive the transition, or 2 optical comb teeth from different combs (beams) interfere with the same effect.

This physical distinction is made in the equations through the polarization vectors $\hat{\varepsilon}_1$ and $\hat{\varepsilon}_2$ for each beam used, and including a provisional absolute relative frequency shift between them. In any case, the associated 2-photon Rabi frequency will be a sum of *all pairs* of optical comb teeth that are spaced $m = \lfloor \omega_{\text{tr}}/\omega_{\text{rep}} \rfloor$ apart, and an additional sum over all states s that mediate that transition and get adiabatically eliminated:

$$\Omega_{m,ij} = \sum_s \sum_{l=-l_0}^{l_0} \frac{g_{l,is} g_{l+m,si}}{2\Delta_{sl}} \tag{A.26}$$

The detunings Δ_{sl} in principle change with both the tooth l and the mediating state s. However, given our pulsed laser, the index l will make Δ_{sl} span a range about 4 times the optical comb's bandwidth $4f_{BW}^{(\text{OC})} = 4/\tau \approx 0.28$ THz:

$$\Delta_{sl_{\max}} - \Delta_{sl_{\min}} \approx 0.28 \text{ THz}$$

which is very small ($< 1\%$) compared to the average values of $\Delta_{sl} = 33$ or 66 THz (depending on which P state mediates the Raman transition). Therefore, the index l will be omitted in (A.26): $\Delta_{sl} \to \Delta_s$. Additionally, the range $\pm l_0$ for the optical comb teeth will be taken to be $\pm\infty$ since the function in the sum decays exponentially with l and this substitution will not have any significant numerical impact, while it simplifies expressions:

$$\Omega_{m,ij} = \sum_s \sum_{l=-\infty}^{\infty} \frac{g_{l,is} g_{l+m,si}}{2\Delta_s}$$

Using (A.25), we get

$$\Omega_{m,ij} = \frac{1}{8} g_0^2 \omega_{\text{rep}} \tau \sum_s \left[\frac{\tilde{C}_{is}\tilde{C}_{sj}}{\Delta_s} \sum_{l=-\infty}^{\infty} \operatorname{sech} \frac{l\omega_{\text{rep}}\tau}{2} \operatorname{sech} \frac{(l+m)\omega_{\text{rep}}\tau}{2} \right]$$

and after some manipulation and approximating the infinite sum with an integral,

$$\Omega_{m,ij} = \frac{1}{4} m\omega_{\text{rep}}\tau \operatorname{csch} \frac{m\omega_{\text{rep}}\tau}{2} \sum_s \frac{\tilde{C}_{is}\tilde{C}_{sj} g_{is,0} g_{sj,0}}{\Delta_s}. \tag{A.27}$$

The states s can in principle be any of the 12 states of the $P_{1/2}$ and the $P_{3/2}$ manifolds. Grouping these together and realizing that the detunings Δ_s are $\Delta = 2\pi \cdot 33.9$ THz for the $P_{1/2}$ state and $\Delta - \omega_F = 2\pi \cdot 66.1$ THz for the $P_{3/2}$ state, we can write

$$\Omega_{m,ij} = \frac{1}{4} m\omega_{\text{rep}}\tau \operatorname{csch} \frac{m\omega_{\text{rep}}\tau}{2} \left[\left(\frac{g^2_{P_{1/2},0}}{\Delta} \sum_{\{s\}\in P_{1/2}} \tilde{C}_{is}\tilde{C}_{sj} \right) + \left(\frac{g^2_{P_{3/2},0}}{\Delta - \omega_F} \sum_{\{s\}\in P_{3/2}} \tilde{C}_{is}\tilde{C}_{sj} \right) \right]. \tag{A.28}$$

Equation (A.28) is a general formula. For transitions that could be of interest between states of the S manifold we can substitute the relevant Clebsch-Gordan coefficients $\tilde{C}_{ab}$ and get a simpler relation where a "base 2-photon Rabi frequency" $\Omega_{m,0}$ emerges (I have ignored all hyperfine splittings since they are $\approx 10^{-4}$ relative

to the fine splitting ω_F in the P manifold and the detunings Δ and $\Delta - \omega_F$):

$$\Omega_{m,|0,0\rangle,|1,0\rangle} = \Omega_{m,0}(\varepsilon_1^- \varepsilon_2^- - \varepsilon_1^+ \varepsilon_2^+) \tag{A.29}$$

$$\Omega_{m,|0,0\rangle,|1,\pm1\rangle} = \pm\Omega_{m,0}(\varepsilon_1^\pm \varepsilon_2^0 + \varepsilon_1^0 \varepsilon_2^\pm) \tag{A.30}$$

$$\Omega_{m,|1,0\rangle,|1,\pm1\rangle} = \Omega_{m,0}(\varepsilon_1^0 \varepsilon_2^\mp + \varepsilon_1^\pm \varepsilon_2^0) \tag{A.31}$$

where $\Omega_{m,0}$ is the polarization-independent 2-photon Rabi frequency for a given $m = \lfloor \omega_{\text{tr}}/\omega_{\text{rep}} \rfloor$.

$$\Omega_{m,0} \equiv \frac{1}{12} m\omega_{\text{rep}}\tau \left(\frac{g^2_{P_{1/2},0}}{\Delta} + \frac{g^2_{P_{3/2},0}}{\omega_F - \Delta} \right) \operatorname{csch} \frac{m\omega_{\text{rep}}\tau}{2} = 0.0171 \cdot \sqrt{\bar{I}_1 \bar{I}_2}.$$

In the last equation, I substituted the values (A.1) corresponding to QSim's pulsed laser, leaving the easily measurable intensities $\bar{I}_1$ and $\bar{I}_2$ of the Raman beams at the ion location as the only variables. Their units are SI, and $\Omega_{m,0}$ is in rad/s. For the beam powers and waists shown in Table A.1, $\Omega_{m,0}/(2\pi) \approx 1040$ kHz. It is easily seen that if none of the Raman beams' wavevector $\vec{k}$ has a projection along the local magnetic field, i.e. has zero π component ($\varepsilon_1^0 = \varepsilon_2^0 = 0$), no Rabi flopping to the Zeeman states will be driven. Given (A.29)–(A.31), we can write for the 2-photon Rabi frequencies btween states of the S manifold:

$$\Omega_{m,ij} = \Omega_{m,0} P_{ij} \tag{A.32}$$

where P_{ij} is a polarization-dependent factor ranging from 0 to 1; for our typical lin⊥lin arrangement in QSim, it is 1. To isolate the dependence on m, which will be useful in a bit, I will also write (A.32) as

$$\Omega_{m,ij} = \Omega_0 P_{ij} \frac{m\omega_{\text{rep}}\tau}{2} \operatorname{csch} \frac{m\omega_{\text{rep}}\tau}{2} \tag{A.33}$$

where I set the base 2-photon Rabi frequency Ω_0 as

$$\Omega_0 \equiv \frac{1}{6}\left(\frac{g^2_{P_{1/2},0}}{\Delta} + \frac{g^2_{P_{3/2},0}}{\omega_F - \Delta}\right)$$

A.5 4thorder Stark shift

Now that we have the 2-photon Rabi frequencies and the detunings for each RF comb tooth, we are ready to write a general 4thorder Stark shift formula, using the general prescription (2.9). For the 4th order Stark shift on state $|i\rangle$ from coupling to state $|j\rangle$ via the k-th RF comb tooth, we can write

$$\Delta E^{(4)}_{k,ij} = \frac{\Omega^2_{k,ij}}{4\delta_{k,ij}}$$

2-photon Rabi frequency between states i and j, if the frequency of RF tooth k was on-resonance and its intensity unchanged

detuning of RF tooth k from the transition $|i\rangle \leftrightarrow |j\rangle$

j: the "coupling" state

i: the state whose frequency is shifted

k: the RF tooth coupling these states

Using (A.6) and (A.33), we get

Single-beam 4-photon Stark shift

$$\Delta E^{(4)}_{k,ij} = \frac{\Omega_0^2}{4} P_{ij}^2 \left(\frac{k\omega_{\text{rep}}\tau}{2}\right)^2 \text{csch}^2 \frac{k\omega_{\text{rep}}\tau}{2} \cdot \frac{1}{k\omega_{\text{rep}} - \omega_{ij}} \tag{A.34}$$

Cross-beam 4-photon Stark shift

$$\Delta E^{(4)}_{k,ij} = \frac{\Omega_0^2}{4} P_{ij}^2 \left(\frac{k\omega_{\text{rep}}\tau}{2}\right)^2 \text{csch}^2 \frac{k\omega_{\text{rep}}\tau}{2} \cdot \left(\frac{1}{k\omega_{\text{rep}} - \omega_{ij} + \delta\omega_{\text{comb}}} + \right. \\ \left. + \frac{1}{k\omega_{\text{rep}} - \omega_{ij} - \delta\omega_{\text{comb}}}\right) \tag{A.35}$$

To find the total 4th order Stark shift on state $|i\rangle$, we thus need to sum over all coupling states $|j\rangle$ and RF comb teeth k[2]. For example, for cross-beam Stark shifts we have:

$$\Delta E^{(4)}_{\text{single},i} = \sum_j \sum_{k \neq m} \Delta E^{(4)}_{k,ij} \tag{A.36}$$

$$= \frac{\Omega_0^2}{4} \sum_j P_{ij}^2 \sum_k \left(\frac{k\omega_{\text{rep}}\tau}{2}\right)^2 \text{csch}^2 \frac{k\omega_{\text{rep}}\tau}{2} \cdot \\ \cdot \left(\frac{1}{k\omega_{\text{rep}} - \omega_{ij} + \delta\omega_{\text{comb}}} + \frac{1}{k\omega_{\text{rep}} - \omega_{ij} - \delta\omega_{\text{comb}}}\right) \tag{A.37}$$

For our usual lin⊥lin Raman beam geometry, the qubit states $|\downarrow\rangle_z$ and $|\uparrow\rangle_z$ can only Raman-couple to one another, therefore the sum over j only has one term, the polarization factor P_{ij} is 1, and $\omega_{ij} = \omega_{\text{HF}}$. For a given pulsed laser (given pulse duration and repetition rate), we can simplify the Stark shift expressions by defining

[2] In the case that an RF tooth k' is resonant with the transition ω_{ij}, it is omitted from the sum.

a *comb factor* for single- and cross-beam geometries:

$$K_{\text{cross},ij}(\delta\omega_{\text{comb}}) \equiv \sum_k \left(\frac{k\omega_{\text{rep}}\tau}{2}\right)^2 \operatorname{csch}^2 \frac{k\omega_{\text{rep}}\tau}{2} \cdot \left(\frac{1}{k\omega_{\text{rep}} - \omega_{ij} + \delta\omega_{\text{comb}}} + \frac{1}{k\omega_{\text{rep}} - \omega_{ij} - \delta\omega_{\text{comb}}}\right) \tag{A.38}$$

$$K_{\text{single},ij} \equiv \sum_k \left(\frac{k\omega_{\text{rep}}\tau}{2}\right)^2 \operatorname{csch}^2 \frac{k\omega_{\text{rep}}\tau}{2} \cdot \frac{1}{k\omega_{\text{rep}} - \omega_{ij}} \tag{A.39}$$

Then, (A.36)–(A.37) become

$$\text{Single-beam:} \quad \Delta E^{(4)}_{\text{single},i} = \frac{\Omega_0^2}{4} \sum_j P_{ij}^2 K_{\text{single},ij} \tag{A.40}$$

$$\text{Cross-beam:} \quad \Delta E^{(4)}_{\text{cross},i} = \frac{\Omega_0^2}{4} \sum_j P_{ij}^2 K_{\text{cross},ij}(\delta\omega_{\text{comb}}). \tag{A.41}$$

where the dependence of the 4th order Stark shift on physical factors is nicely separated: Ω_0 codifies the light intensity, P_{ij} codifies the polarization, and the comb factors K codify the frequency dependence. All these properties of the light pertain to the ions' location. To make this statement explicit, I will write for both single- and cross-beam Stark shifts

$$\Delta E_i^{(4)} = \zeta \bar{I}_1 \bar{I}_2 \sum_j P_{ij}^2 K_{ij}$$

where ζ does not depend on any light properties. The light intensities $\bar{I}_1$ and $\bar{I}_2$ correspond to the two optical combs interfering. If that's a single-beam Stark shift, we replace their product with $\bar{I}_1^2$. This expression will be very useful when translating light intensity noise to Stark shift noise.

Appendix B: The state discrimination algorithm

This section aims to briefly present the state discrimination algorithm used in QSim, which is the last stage of qubit readout. Is is often viewed as a trivial stage of quantum simulation, but it often comes up in data analysis, and a quick presentation will be hopefully useful.

Qubit readout (Fig. B.1) is a multi-stage process. First, the ions are illuminated with a resonant 369 nm "detection" beam. For details, see [13] or [27]. When an ion interacts with its first detection photon, its spin wavefunction is projected to either $|\downarrow\rangle_z$ or $|\uparrow\rangle_z$. From there on, it either undergoes the detection cycling transition absorbing and reemitting photons, or it does not emit at all. Second, the imaging system collects a small part of this fluorescence and focuses it on the sensor of an EMCCD camera (Andor iXon Ultra 897). The camera system converts these photons to current, and ultimately communicates a two-dimensional array to a PC, containing the counts for every pixel of the rectangular acquisition area. The acquisition rate of the camera is currently the bottleneck in the repetition rate of our experimental sequences, because after every exposure it has to undergo a "keep clean" cycle that cleans the sensor from accumulated charge, even if it was just read out. Finally, this 2D array of numbers is input to an algorithm that determines

whether an ion was bright or dark, i.e. $|\uparrow\rangle_z$ or $|\downarrow\rangle_z$, for every ion individually.

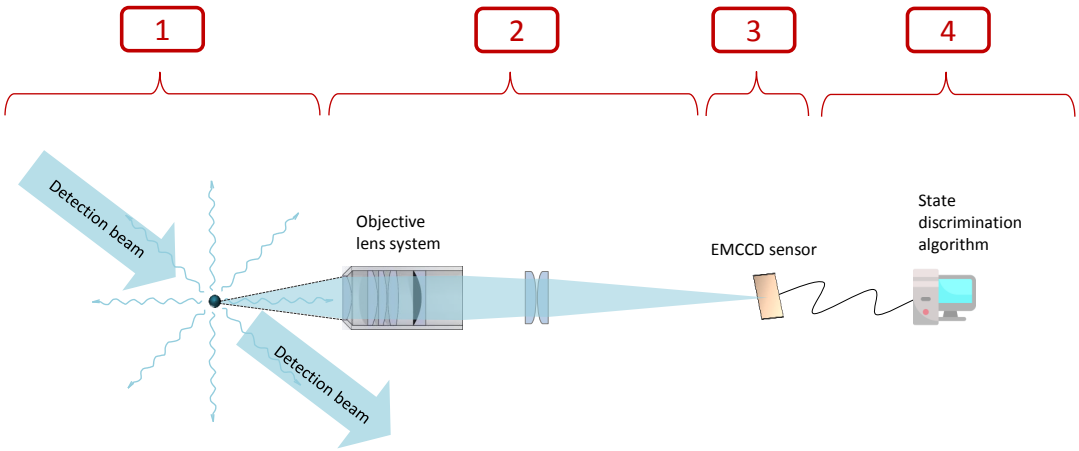

Figure B.1: **The flow of qubit state readout. 1.** The 369 nm detection light causes the ions to fluoresce or not, depending on their state. **2**. The imaging system collects a fraction of this fluorescence and focuses it to the EMCCD sensor, with an overall magnification of about 70×. **3**. The EMCCD sensor's pixels convert incident photons to current, which is converted to a digital signal indicating the number of electron counts by the Andor electronics. **4**. The QSim computer algorithm decides if these counts indicate an ion in the bright or the dark state.

Fig. B.2 showcases the flow of the calibration procedure for state discrimination. This procedures first identifies which subsets of pixels (ROI, region of interest) within the large exposure area are encircling each ion. After finding these ROIs, it identifies the optimal fluorescence threshold to classify a certain ROI as bright or dark. One would naively think that any collected 369 nm fluorescence at all signifies a bright ion. This is not true, since we have a non-negligible amount of 369 nm light directed into the imaging optics, that is not ion fluorescence but coming straight from the laser. Adding to that, camera pixels have a finite probability to "fire" and appear bright even at the absence of a photon; finally, there is a small but finite chance that a dark ion will be off-resonantly excited by the 369 nm light and start

fluorescing (off-resonant pumping).

To start, the ion chain is illuminated with Doppler cooling light and the camera is exposed for a long duration (of the order of 20 ms). This long duration allows for a large amount of ion fluorescence to accumulate on the sensor, facilitating drawing the ROI around each of the bright disks (ions). Once these ROIs are found, the ion chain is initialized to the bright state and illuminated with resonant detection light for the normal exposure time, of the order of 0.3 ms. These bright state shots are repeated 1000 times. The image acquired (stage d in Fig. B.2) is significantly dimmer than the first Doppler-cooling shot. After that, the chain is initialized in the dark state and the camera is exposed again 1000 times; any counts registered now are undesired noise. A procedure is then followed, consisting of comparing a property of the bright pictures for each ion to the same property calculated for the dark ones for that ion, and finding a value that optimally delineates between these two.

Initially we were using as this measure the natural choice of the total counts enclosed in the ion ROI. Recently, however, we switched to a slightly more elaborate measure because it was yielding a higher degree of separation between the dark and the bright distributions (by $0.1 - 0.5\%$, which matters for fidelities that are $\approx 97\%$ to begin with). For this measure, we average the counts for the brightest l-th and $(l+1)$-th pixels, for example the 4th and 5th brightest ones. The dark and bright distributions of this measure (quantiles) are slightly more separated than just the total ROI counts. The reason for that is conjectured to be its robustness against camera noise and pixels firing randomly: Not more than 1-3 pixels per ROI are

expected to have this erroneous behavior. Therefore, by discarding the brightest say 3 pixels we discard all of these spurious bright events. Of course, in the case of a bright ion, we are also throwing away pixels that registered legitimate fluorescence counts. However, a fluorescing ion will cause more than 4 pixels in its ROI to register high counts, with high certainty. Thus, we can afford to throw away some of them and still be left with some bright pixels if the ion was bright.

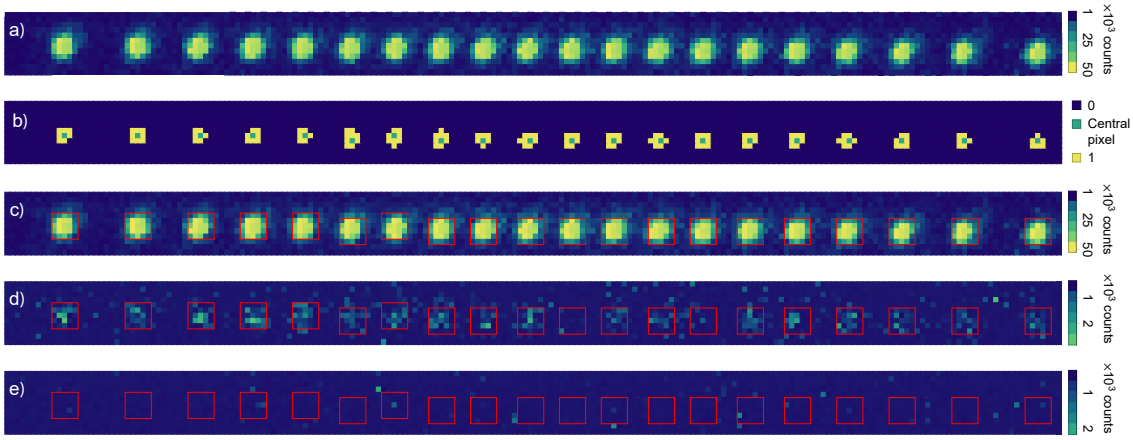

Figure B.2: **Stages of the state discrimination algorithm.**
a) The cooling shot (i.e. the array of pixel counts) is loaded into memory.
b) The raw cooling shot is rescaled so that the counts are expressed in a range [0-1] for ease, then smoothened using a gaussian filter of order the ion's radius as it appears on the camera, and importantly, converted to a binary landscape by identifying "sharp edges" with the `MaxDetect` command of Wolfram Mathematica®. At the end of this stage we identify the geometrical center for each of the yellow "islands" and round its {x,y} coordinates to integers; this location will be the identified center for this ion (green pixel).
c) The rectangular ROI rectangles are calculated for each ion. Typical choices for their size are 5×5 and 7×7, but they can be tailored to account for smaller ion spacing at the center of the chain, requiring 5×7 ROI dimensions for these ions.
d) A typical bright shot, with the EMCCD set to maximum EM gain.
e) A typical dark shot (same parameters as the bright shot).

Appendix C: The data takers

Humility.

The data takers patiently gather the crops

from their electric plantations.

Pushing buttons and aligning optics,

never cursing at the slip of a tool

or at the loss of a day.

Never arrogant,

with thoughts of a distant warm future

drawing a faint smile on their lips.

Gleams on a monitor is what they are looking for.

And they do not quit until their baskets are heavy

with the precious fruits of their labor.

Oh, how I wish it was easier to be like them.

Bibliography

[1] I. M. Georgescu, S. Ashhab, and Franco Nori. Quantum simulation. *Reviews of Modern Physics*, 86, 2014.

[2] C. Monroe, W. C. Campbell, L. M. Duan, Z. X. Gong, A. V. Gorshkov, P. Hess, R. Islam, K. Kim, N. Linke, G. Pagano, P. Richerme, C. Senko, and N. Y. Yao. Programmable quantum simulations of spin systems with trapped ions. *arXiv:1912.07845v2 [quant-ph]*, 2020.

[3] P. W. Hess, P. Becker, H. B. Kaplan, A. Kyprianidis, A. C. Lee, B. Neyenhuis, G. Pagano, P. Richerme, C. Senko, J. Smith, W. L. Tan, J. Zhang, and C. Monroe. Non-thermalization in trapped atomic ion spin chains. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 375, 2017.

[4] P. Richerme, C. Senko, J. Smith, A. Lee, S. Korenblit, and C. Monroe. Experimental performance of a quantum simulator: Optimizing adiabatic evolution and identifying many-body ground states. *Physical Review A*, 88, 2013.

[5] Philip Richerme, Zhe Xuan Gong, Aaron Lee, Crystal Senko, Jacob Smith, Michael Foss-Feig, Spyridon Michalakis, Alexey V. Gorshkov, and Christopher Monroe. Non-local propagation of correlations in quantum systems with long-range interactions. *Nature*, 511:198–201, 2014.

[6] J. Smith, A. Lee, P. Richerme, B. Neyenhuis, P. W. Hess, P. Hauke, M. Heyl, D. A. Huse, and C. Monroe. Many-body localization in a quantum simulator with programmable random disorder. *Nature Physics*, 12:907–911, 10 2016.

[7] Brian Neyenhuis, Jiehang Zhang, Paul W Hess, Jacob Smith, Aaron C Lee, Phil Richerme, Zhe-Xuan Gong, Alexey V Gorshkov, and Christopher Monroe. Observation of prethermalization in long-range interacting spin chains, 2017.

[8] J. Zhang, P. W. Hess, A. Kyprianidis, P. Becker, A. Lee, J. Smith, G. Pagano, I.-D. Potirniche, A. C. Potter, A. Vishwanath, N. Y. Yao, and C. Monroe. Observation of a discrete time crystal. *Nature*, 543, 2017.

[9] J. Zhang, G. Pagano, P. W. Hess, A. Kyprianidis, P. Becker, H. Kaplan, A. V. Gorshkov, Z. X. Gong, and C. Monroe. Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator. *Nature*, 551:601–604, 2017.

[10] S. Olmschenk, K. C. Younge, D. L. Moehring, D. N. Matsukevich, P. Maunz, and C. Monroe. Manipulation and detection of a trapped yb+ hyperfine qubit. *Physical Review A - Atomic, Molecular, and Optical Physics*, 76, 2007.

[11] Steven M. Olmschenk. Quantum teleportation between distant matter qubits (phd thesis), 2009.

[12] R W Berends, E H Pinnington, B Guo, and Q Ji. Beam-laser lifetime measurements for four resonance levels of yb 11, 1993.

[13] M. Acton, K.-A Brickman, P C Haljan, P J Lee, L Deslauriers, and C Monroe. Near-perfect simultaneous measurement of a qubit register. *arXiv:quant-ph/0511257v2*, 2006.

[14] Rachel Noek, Geert Vrijsen, Daniel Gaultney, Emily Mount, Taehyun Kim, Peter Maunz, and Jungsang Kim. High speed, high fidelity detection of an atomic hyperfine qubit. *Optics Letters*, 38:4735, 2013.

[15] A. H. Myerson, D. J. Szwer, S. C. Webster, D. T.C. Allcock, M. J. Curtis, G. Imreh, J. A. Sherman, D. N. Stacey, A. M. Steane, and D. M. Lucas. High-fidelity readout of trapped-ion qubits. *Physical Review Letters*, 100, 2008.

[16] A. H. Burrell, D. J. Szwer, S. C. Webster, and D. M. Lucas. Scalable simultaneous multiqubit readout with 99.99fidelity. *Physical Review A - Atomic, Molecular, and Optical Physics*, 81, 2010.

[17] Kenneth R. Brown, Aram W. Harrow, and Isaac L. Chuang. Arbitrarily accurate composite pulse sequences. *Physical Review A - Atomic, Molecular, and Optical Physics*, 70, 2004.

[18] K. Kim, M. S. Chang, R. Islam, S. Korenblit, L. M. Duan, and C. Monroe. Entanglement and tunable spin-spin couplings between trapped ions using multiple transverse modes. *Physical Review Letters*, 103, 2009.

[19] Wolfgang Paul. Electromagnetic traps for charged and neutral particles. *Reviews of Modern Physics*, 62, 1990.

[20] R. Islam, E. E. Edwards, K. Kim, S. Korenblit, C. Noh, H. Carmichael, G. D. Lin, L. M. Duan, C. C. Joseph Wang, J. K. Freericks, and C. Monroe. Onset of a quantum phase transition with a trapped ion quantum simulator. *Nature Communications*, 2, 2011.

[21] D.J. Wineland, C. Monroe, W.M. Itano, D. Leibfried, B.E. King, and D.M. Meekhof. Experimental issues in coherent quantum-state manipulation of trapped atomic ions. *Journal of Research of the National Institute of Standards and Technology*, 103, 1998.

[22] J. Earnshaw. On the nature of the molecular forces which regulate the constitution of the luminiferous ether. *Trans. Camb. Phil. Soc.*, 7:97–112, 1839.

[23] H.G. Dehmelt. Radiofrequency spectroscopy of stored ions i: Storage.

[24] Klaus Mølmer and Anders Sørensen. Multiparticle entanglement of hot trapped ions. *Physical Review Letters*, 82, 1999.

[25] Shi Liang Zhu, C. Monroe, and L. M. Duan. Trapped ion quantum computation with transverse phonon modes. *Physical Review Letters*, 97, 2006.

[26] Kazi Rajibul Islam. Quantum simulation of interacting spin models with trapped ions (phd thesis), 2012.

[27] Crystal Rosalie Senko. Dynamics and excited states of quantum many-body spin chains with trapped ions (phd thesis), 2014.

[28] Aaron C. Lee. Engineering a quantum many-body hamiltonian with trapped ions (phd thesis), 2016.

[29] A. C. Lee, J. Smith, P. Richerme, B. Neyenhuis, P. W. Hess, J. Zhang, and C. Monroe. Engineering large stark shifts for control of individual clock state qubits. *Physical Review A*, 94, 2016.

[30] K. Kim, S. Korenblit, R. Islam, E. E. Edwards, M. S. Chang, C. Noh, H. Carmichael, G. D. Lin, L. M. Duan, C. C. Joseph Wang, J. K. Freericks, and C. Monroe. Quantum simulation of the transverse ising model with trapped ions. *New Journal of Physics*, 13, 2011.

[31] Jacom Smith. Quantum thermalization and localization in a trapped ion quantum simulator (phd thesis), 2016.

[32] C.-C. Joseph Wang and J. K. Freericks. Intrinsic phonon effects on analog quantum simulators with ultracold trapped ions. *Physical Review A*, 86, 2012.

[33] D F James and J Jerke. Effective hamiltonian theory and its applications in quantum information. *Can. J. Phys.*, 85, 2007.

[34] H. J. Metcalf and P. van der Straten. Laser cooling and trapping of atoms. *J. Opt. Soc. Am. B*, 20, 2003.

[35] Yukai Wu. Effects of phonon excitation in the ion trap simulation (private communication), 2017.

[36] Norman Y. Yao and Chetan Nayak. Time crystals in periodically driven systems. *Physics Today*, 71, 2018.

[37] A. Kyprianidis, F. Machado, W. Morong, P. Becker, K. S. Collins, D. V. Else, L. Feng, P. W. Hess, C. Nayak, G. Pagano, N. Y. Yao, and C. Monroe. Observation of a prethermal discrete time crystal. *arXiv:2102.01695 [quant-ph]*, 2021.

[38] V. Khemani, R. Moessner, and S. L. Sondhi. A brief history of time crystals. *arXiv preprint arXiv:1910.10745*, 2019.

[39] Francisco Machado, Dominic V. Else, Gregory D. Kahanamoku-Meyer, Chetan Nayak, and Norman Y. Yao. Long-range prethermal phases of nonequilibrium matter. *Physical Review X*, 10, 2020.

[40] Dominic V. Else, Bela Bauer, and Chetan Nayak. Prethermal phases of matter protected by time-translation symmetry. *Physical Review X*, 7, 2017.

[41] Wineland D. J., C. Monroe, W. M. Itano, B. E. King, D. Leibfried, D. M. Meekhof, C. Myatt, and C. Wood. Experimental primer on the trapped ion quantum computer. *Fortschr. Phys.*, 46:363–390, 1998.

[42] T. A. Savard, K. M. O'Hara, and J. E. Thomas. Laser-noise-induced heating in far-off resonance optical traps. *Physical Review A*, 56, 8 1997.

[43] Q. A. Turchette, Kielpinski, B. E. King, D. Leibfried, D. M. Meekhof, C. J. Myatt, M. A. Rowe, C. A. Sackett, C. S. Wood, W. M. Itano, C. Monroe, and D. J. Wineland. Heating of trapped ions from the quantum ground state. *Physical Review A*, 61, 2000.

[44] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland. Quantum dynamics of single trapped ions. *Reviews of Modern Physics*, 75, 3 2003.

[45] G Pagano, P W Hess, H B Kaplan, W L Tan, P Richerme, P Becker, A Kyprianidis, J Zhang, E Birckelbaw, M R Hernandez, Y Wu, and C Monroe. Cryogenic trapped-ion system for large scale quantum simulation. *Quantum Science and Technology*, 4, 10 2018.

[46] J. H. Wesenberg. Electrostatics of surface-electrode ion traps. *Physical Review A*, 78, 2008.

[47] Wade Hodson and Christopher Jarzynski. Energy diffusion coefficient for a particle in an anharmonic trap, subject to weak noise (writeup and private communication), 2020.

[48] E. Paladino, Y. M. Galperin, G. Falci, and B. L. Altshuler. 1—f noise: Implications for solid-state quantum information. *Reviews of Modern Physics*, 86, 4 2014.

[49] F.N. Hooge and P.A. Bobbert. On the correlation function of 1—f noise. *Physica B: Condensed Matter*, 239, 1997.

[50] D. Hayes, D. N. Matsukevich, P. Maunz, D. Hucul, Q. Quraishi, S. Olmschenk, W. Campbell, J. Mizrahi, C. Senko, and C. Monroe. Entanglement of atomic qubits using an optical frequency comb. *Physical Review Letters*, 104, 2010.

[51] Jonathan Mizrahi. Ultrafast control of spin and motion in trapped ions (phd thesis), 2013.

[52] C. J. Foot. *Atomic Physics*. Oxford University Press, 2005.

[53] E. H. Pinnington, G. Rieger, and J. A. Kernahan. Beam-laser measurements of the lifetimes of the 6p levels in yb ii. *Physical Review A*, 56, 1997.

[54] S. Olmschenk, D. Hayes, D. N. Matsukevich, P. Maunz, D. L. Moehring, K. C. Younge, and C. Monroe. Measurement of the lifetime of the 6p2p1/2 level of yb+. *Physical Review A*, 80, 2009.