MASTERARBEIT | MASTER'S THESIS

Titel | Title

Extending Free Expansion Times of Charged Levitated Nanoparticles in Vacuum

verfasst von | submitted by

David Steiner BSc

angestrebter akademischer Grad | in partial fulfilment of the requirements for the degree of

Master of Science (MSc)

Wien | Vienna, 2026

Studienkennzahl lt. Studienblatt | Degree programme code as it appears on the student record sheet:	UA 066 876
Studienrichtung lt. Studienblatt | Degree programme as it appears on the student record sheet:	Masterstudium Physics
Betreut von | Supervisor:	Assoz. Prof. Dr. Thomas Nikolai Kiesel
Mitbetreut von | Co-Supervisor:	Dr. Yaakov Fein BA MA MSc

Acknowledgements

I would like to express my heartfelt gratitude to all the people, from Vienna and beyond, who were along for this journey. Without their continuous encouragement, support, and help, my studies might never have been completed – or certainly not in a reasonable timeframe.

Let me first thank the supervisors and group leaders, Yaakov Y. Fein, Nikolai Kiesel, and Markus Aspelmeyer. I am very grateful for the opportunity to work in such a great group, and for the support I received from all of you. Whether it was feedback on posters, abstracts, recommendation letters, or ideas on improving this thesis – you always took the time and offered great suggestions on my first drafts.

From day one, it felt like I was part of a team and truly supported. I appreciate all the discussions, be it in group meetings or in chats, that have heavily shaped this thesis. For this I also want to thank the whole group, with a specific thanks going to the 3rd floor office regulars and the park lunchers.

Further thanks go to my family, for enabling me to move to a different city and continuously supporting me throughout my studies. Whether it was late-night pickups when I came to visit, food and coffee beans I would come back to Vienna with, or the odd tool you brought to fix things in my apartment – you were always there for me and I am incredibly grateful for that.

Studies were not ompletely made up of work, but also of fun. I want to acknowledge the people who stuck around since day one as well as those who joined along the way. Whether it was working on coursework or attending a fun gathering, you were there, and that guaranteed it was worth the time. With a heavy heart that some of you have turned your backs on Vienna, I am happy about the time we spent together and that we still manage to keep in touch.

Lastly, I want to thank Flora. From kindling the thought of studying natural science in the first place, to just recently giving feedback on this very thesis – the continuous moral support, thought-provoking and insightful discussions, and your general presence in my life is something I am grateful for every second. You are a pleasure to be around, brilliantly smart, and amazing.

AI Notice: For the work presented in this thesis, some amount of assistance from AI tools (e.g. Copilot, autocomplete) was used for parts of the coding tasks to produce the figures in this document. Additionally, spellcheck and grammar tools which also contain language model and AI components (e.g. DeepL, Claude) were also used.

Abstract

Coherent expansion of a quantum wavepacket is a necessary fundamental step in many experimental protocols to prepare, manipulate, and read out non-classical states, especially for levitated dielectric nanoparticles where the ground state extent is on the order of 10 pm. Free evolution is potentially the most straightforward technique to achieve this state expansion. Free evolution means releasing the particle from its harmonic trapping potential, allowing its position variance to grow, expanding the state. While it is both a simple and effective expansion scheme, it does require compensation of static forces to avoid mean displacement of the wavepacket during the free evolution.

Using an optical trap-release-recapture sequence involving a charged nanoparticle we developed methods to measure and counteract static forces. Compensating gravity and stray electric fields, in three dimensions, we demonstrate free evolution for 1 ms with subsequent recapture. To achieve this, we additionally developed methods to map and correct for force cross-talk between axes, as well as to control in a deterministic manner the environmental charge state.
Coupled with a short harmonic pulse in between two free expansions, we have performed recompression and time-of-flight protocols, speaking for the coherence of our approach as well as its usefulness to further generate squeezed mechanical states.

Kurzfassung

Die kohärente Ausdehnung eines quantenmechanischen Wellenpakets ist ein notwendiger, fundamentaler Schritt in vielen experimentellen Protokollen, um nicht-klassische Zustände zu erzeugen, zu manipulieren und diese auch auszulesen, insbesondere für optisch levitierte, dielektrische Nanoteilchen, bei denen die Ausdehnung des Grundzustands im Potential in der Größenordnung von 10 pm liegt. Freie Evolution ist wahrscheinlich die naheliegendste Technik, um diese Zustandsausdehnung zu erreichen. Freie Evolution bedeutet in diesem Fall, das Teilchen aus seinem harmonischen Fallenpotential freizulassen, wodurch sich die Positionsvarianz vergrößert und sich der Zustand koheränt ausdehnt. Obwohl es sich um eine einfache und effektive Technik zur Ausdehnung handelt, erfordert eben genau jenes Ausschalten des Fallenpotentials eine Kompensation statischer Kräfte, um zu verhindern, dass sich das Teilchen während der freien Evolution aus dem Fallenpotential bewegt.

Mit wiederholtem Fangen im optischen Potential, Freilassen und Wiederauffangen eines geladenen Nanoteilchens wurden Methoden zur Messung und Kompensation statischer Kräfte entwickelt. Die Kompensation von Gravitation und elektrostatischen Feldern in drei Dimensionen ermöglicht infolgedessen eine freie Evolution von 1 ms mit anschließendem Wiederauffangen des Nanoteilchens. Um dies zu erreichen, haben wir zusätzlich Methoden entwickelt, den *cross-talk* zwischen den Elektrodenachsen zu messen und zu korrigieren. Ebenso wurden Wege entwickelt, die Ladungsverteilung der Umgebung kontrolliert zu manipulieren.

Wird nun noch ein kurzer harmonischer Puls zwischen zwei freien Expansionen hinzugefügt, so können wir Rekompressions- und Time-of-Flight-Protokolle durchführen, welche die Kohärenz des Prozesses sowie die Anwendbarkeit zur weiteren Erzeugung von mechanischen *Squeezed States* zeigen.

Contents

List of Figures

Notes on the publishing of part of this thesis

Some parts of this thesis have been published in the following peer-reviewed journal article:

> D. Steiner, Y. Y. Fein, G. Meier, S. Lindner, P. Juschitz, M. A. Ciampini, M. Aspelmeyer and N. Kiesel, 'Free expansion of a charged nanoparticle via electrostatic compensation', *Appl. Phys. Lett.*, vol. 127, no. 19, p. 191 103, 10th Nov. 2025, ISSN: 0003-6951. DOI: `10.1063/5.0303460`. Accessed: 29th Nov. 2025. [Online]. Available: `https://doi.org/10.1063/5.0303460`.

Nevertheless, this thesis is an independent work and contains many points and data that are not part of the publication. All plots and figures were made specifically for this thesis, even if the underlying data was also used in the publication, unless otherwise noted. No text from the publication has been copied verbatim into this thesis, but at points the structure naturally lends itself to similar phrasing.

1. Introduction

1.1. Motivation

Optically trapped dielectric nanoparticles in high vacuum have emerged as a promising platform for exploring macroscopic quantum physics due to their inherent isolation from the environment [2, 3] and their excellent controlability. Unlike traditional clamped resonators, these levitated objects lack a physical attachment to the environment. This isolation and lack of clamping losses allow the mechanical quality factor of the centre-of-mass motion to reach values above 10^{10} at high vacuum levels [2], comparable to the performance of cryogenic clamped systems, such as those used for single-electron spin detection [4], but achieved without the overhead of refrigeration. Historically, such transducers have been at the forefront of experimental physics, enabling breakthrough measurements ranging from the imaging of individual covalent bonds in atomic force microscopy [5] to the landmark detection of gravitational waves [6]. Additionally, optically levitating nanoparticles offers outstanding control over their motional degrees of freedom. Consequently, these levitated systems have already demonstrated attonewton-scale force sensitivities [7], establishing them as exceptional transducers. By leveraging these high quality factors and the resulting long coherence times, researchers have achieved unprecedented control over the levitated particle's dynamics [8], including cooling the centre-of-mass motion to the quantum ground state [9, 10, 11, 12]. This level of control transitions levitated systems from purely classical sensors into the quantum regime, where they join a large group of high-precision mechanical transducers.

A common challenge in the field pertains to the measurement of the minute perturbations acting on the mechanical transducer, limited by the quantum nature of the electromagnetic probe field used to read out the transducer position due to measurement backaction and imprecision [13]. This phenomenon is referred to as the standard quantum limit. It is evident that within the field of physics, there is a strong inclination to establish limits and subsequently seek methods to surpass these same boundaries. In response to this standard quantum limit, novel techniques were developed to overcome it by exploiting quantum correlations between the amplitude and phase quadrature of the probe after it interacted with the transducer, by using squeezed probe light with engineered quantum correlations. Recently, instead of focusing on the probe field, quantum correlations of the transducer's motion have been leveraged for sensing impulsive forces, requiring a squeezed transducer state. One way these schemes can benefit from force-compensated free evolutions is to achieve such a squeezed state by using evolutions of the mechanical transducer. For readout, one can also use free evolution as a time of flight measurement [14, 15, 16].

Pushing the boundary of what can be controlled quantum-mechanically does not only improve sensitivity, it also probes the regime where quantum mechanics itself has not yet been tested. Quantum mechanics, despite being one of the most widely tested theories to date, has not been verified for objects large enough to be considered macroscopic, and it is precisely the mass scale targeted by levitated-particle experiments that falls into this uncharted territory. The goal therefore is to now keep the excellent properties of optically levitated nanoparticles as tiny sensors, while pushing them back up into a macroscopic regime where quantum mechanics has not been tested yet. One robust way to prove a system's quantum nature is to conduct interference experiments.

Until now, the most macroscopic interfered isolated single objects were sodium clusters with masses up to $170\,000\,\mathrm{Da} \approx 2.83 \times 10^{-19}\,\mathrm{g}$ [17] and a diameter of approximately 8 nm. These experiments were performed in a typical matter wave interferometer setup, where the sodium is sent through a series of gratings to create and measure the interference pattern at the end. One fundamental limit to these types of setups is however the very fact which makes them work: the key figure is the de Broglie wavelength of the particle, which shrinks with increasing mass of the particle and requires slower and slower particles to achieve the same interference pattern. This quickly leads to impractical setups, as the time of flight for heavy particles becomes long due to the required low velocities. This requires extremely long, extremely good vacuum chambers and control over external perturbations acting on the particle during this time. Additionally, as the time of flight increases, the particle becomes more and more displaced due to gravity.

One way to not face these particular challenges whilst taking another leap in mass of orders of magnitude is to use levitated particles, which were already proven to be excellent sensors. The idea is to leverage the knowledge gained from these smaller systems to preserve the desired properties — such as quality factors and control — while pushing the mass further up. With such ideas in mind, protocols have been proposed to create and measure macroscopic quantum states. These protocols rely heavily on the ability to increase the spatial extent of the particle's wavefunction, using interference patterns measured at the end of the protocol run as verification of the quantum nature of the state [18] [19]. These interference patterns would show up even though the particle remains in situ, ideally not moving away from the trap position during the whole run.

These heavy particles could then in turn still be used to also sense external perturbations, as the large spatial extent of the wavefunction would lead to an enhanced sensitivity to forces acting on the particle during the protocol. This could allow for tests of collapse models [20], which predict a mass-dependent breakdown of quantum mechanics, as well as searches for new fundamental interactions [21, 22] or dark matter.

As the particle gets heavier, the spatial extent of the wavefunction gets smaller, which makes the implementation of readout and state manipulation pulses more difficult. To remedy this, we need to develop a method to increase the state size, which is the motivation for this thesis.

1.2. Overview of a levitated matter wave interferometry protocol

To make clear what constrains the free evolution and in what time scales we are interested in as a primary goalpost, we will outline one of the proposed protocols [18], which is being pursued currently in the group.
The protocol consists of five steps (see [18, Fig. 1]) but since the exact timings will likely be modified as the protocol is further put into practice, we will just give an overview of the steps and the order of magnitude of the timescales as proposed in the published version, which are in the microsecond to millisecond range. We will also touch upon why charged particles are a challenge worth pursuing in the first place. The steps are as follows:

0) Initialisation of the ground state in a harmonic potential,

1) free evolution, for a time $\tau_1 \approx 1 - 2\,\mathrm{ms}$,

2) pulsed interaction with a cubic potential, for a time of $\tau_2 \approx 1\,\mathrm{\mu s}$,

3) free evolution, for a time $\tau_3 \approx 1\,\mathrm{ms}$,

4) evolution in an inverted potential, for a time $\tau_4 \approx 100\,\mathrm{\mu s}$.

5) measurement of the particle position in a harmonic potential, which should show an interference pattern if the protocol is successful.

As the pulses are currently envisioned to be implemented via electromagnetic fields, the particle needs to be charged, which means that the protocol relies on the ability to compensate the resulting electrostatic forces during free evolution. The same applies for gravity, which is also compensated by an electric field, since after the last step, the particle position is measured in the original harmonic potential and then re-initialized for the next run. The process therefore relies for increasingly longer evolution times ever more on the ability to keep the particle in the same place throughout the protocol. Only then can the same particle be used for multiple runs, which is necessary to gather enough statistics for the interference pattern measurement at the end. In subsection 1.3.3 we will further discuss the state of the art of free evolution experiments with optically trapped nanoparticles in different setups, while also highlighting why the work of this thesis increases the state of the art by order of magnitude in terms of free evolution times. Discharging the particle and using a neutral particle would of course circumvent the issue of electrostatic force compensation, but then the particle would be displaced by gravity during the free evolution, which needs to be compensated for in a different way, which we will briefly touch upon in the next few paragraphs.

1.3. Methods for state expansion, an overview

The main use for the free evolution in the above protocol is to coherently increase the spatial extent of the particle's wavefunction in step 1) and then use it as a time of flight readout in step 3). As the wavefunctions zero point motion scales as $z_{\text{zp}} = \sqrt{\hbar/2m\Omega_z}$, with m the mass of the particle and Ω_z the trap frequency, the initial state extent is tiny, on the order of tens of picometers for a 100 nm particle. This makes the pulses and measurement in steps 2), 4) and 5) impossible to implement due to noise and the way the measurement is planned to work. By letting the particle evolve freely for a time τ_1 and τ_3, the wavefunction can expand large enough to implement the pulses and measurement. Longer free evolutions therefore generally allow for relaxed constraints on pulses and measurement.

There are multiple ways to achieve the same goal of large state expansion. To illustrate what makes the static force compensation method of this thesis special, we will briefly discuss some of them. This list is by no means meant to be exhaustive and the ideal method depends on the specific experiment/protocol/setup/constraints.

1.3.1. Inverted harmonic expansion

Employing an inverted harmonic potential on which the particle expands on would actively accelerate the state expansion exponentially, making the process faster. This would be beneficial not only because of the increased repetition rate, but it would also be less sensitive to position localization decoherence from stochastic collisions with the residual gas due to the reduced timescale. This dubbed "coherent inflation" could therefore be a way to achieve bigger states with the same vacuum level if one can make sufficiently sure that the state is located exactly at the apex of the inverted harmonic. [23]

Early experiments make it seem unclear if ensemble decoherence over multiple runs due to variations and noise of this perfect positioning is a fundamental issue with the technique in terms of [24]. For inverted harmonic state expansion, the decoherence length currently decays for longer expansion times [24, Fig. 4], presumably due to the fact that two different methods of generating the two potentials are employed, leading to misalignment in their respective extrema between runs as the reason for this decoherence. Progress however is being made where trapping the particle on top of the inverted harmonic potential is achieved [25, 26], maybe also providing a way to not face the issue with distance between trap extrema.

1.3.2. Frequency jumping

The final method not utilising free evolution that we will briefly introduce here is called frequency jumping and is a continuous method for squeezing the state of the particle. The name comes from the idea of quickly changing the trap stiffness (frequency) from a high value to a low value and back again. This sudden change in frequency leads to an increase in position variance, as the particle's wavefunction expands in the lower frequency potential.

When the oscillation frequency is quickly varied between ω_1 to ω_2, the phase space variance is changed by a factor of $\exp(-r)$, where $r = \ln(\omega_1/\omega_2)/2$ is a squeezing parameter. By then timing the duration of the pulses correctly, this method leads to squeezing of the state by a factor of $\exp(-2r)$ [8, 27, 28]. We will later also dive into what "squeezing" means in section 2.2, but for now let us just succinctly introduce it as a way to increase the state extent in one phase space quadrature, while decreasing it in the other quadrature even below the ground state extent.

1.3.3. Free evolution and state of the art

Free evolution of neutral particles

In recent publications [10, 29], free evolutions of 68 µs and 125 µs are performed on a similar sized, but neutral particle as a time of flight measurements to measure the velocity distribution and squeezing of the centre of mass motion.
Another paper focusing mainly on measuring static forces [30] performs free evolutions of up to 100 µs on neutral particles as well.

Free evolution of neutral particles with tracking trap

Circumventing the issue of a neutral particle displacing along the direction of the static gravitational force, for longer free evolutions the trap can be deflected via an acousto-optic modulator (AOM) at different frequencies to follow the particle during the free fall. This allows recapture after 250 µs of free evolution time [31]. The AOM drive frequency at the end of the free fall is then ramped back to the original frequency and therefore trap position to restart the protocol. Still with this method there is another challenge in sight: Once the particle displaces too far from the original position, the AOM cannot easily follow the particle, also as the trapping light then passes through the optical system at a different angle, meaning the optical setup could limit the free fall duration at some point. This could be circumvnted by going up in complexity, for example by catching the particle in a second, different, trap after the free fall, again coming with it's own caveats and challenges.

Free evolution of neutral particles in a free falling frame

In a free falling frame, there are no effects of gravity, allowing for longer free evolution times without the need for active stabilization or complicated movement to the predicted particle position after the free evolution. This can be achieved by placing the entire experimental setup in a drop tower [32] or even a satellite [33].
These methods are costly, require complicated setups highly robust to accelerations and vibrations during the drop or launch, and are not easily accessible for quick iterations of experiments. Furthermore, to the best of our knowledge, no experiments including free evolutions of charged particles in such a free falling frame have been performed up to this date.

1.3.4. Free evolution of charged particles with DC-Compensation

The method of this thesis is to use the usually negative aspect of the charged particles to our advantage, by applying DC voltages to the electrodes to compensate for the resulting electrostatic forces during free evolution but also for gravity. This allows for long free evolutions without the need for any movement of the trap or the entire setup, while still being able to use charged particles, which are required for the pulses in the protocol we outlined above. As gravity cannot be turned off, this method is effectively the only way to compensate it instead of working around it during the free evolution time.

Structure of the Thesis

Since understanding the force-compensation technique and the subsequent measurements requires a solid understanding of how a particle evolves freely and what happens in phase space during that evolution, we will start with the theoretical background in Chapter 2. There we will first cover the physics of optical trapping, deriving the harmonic potential seen by the particle and the equation of motion that governs its dynamics. We will then discuss squeezing and the evolution of the full covariance matrix, as this is the natural language for describing what happens to the state during a free evolution. The Duffing nonlinearity of the trap potential is treated separately, since it becomes relevant at the large motional amplitudes encountered after recapture. We will also introduce the two main figures of merit used throughout the thesis – the mean energy increase and the growth of position variance – and close the chapter with a discussion of the dominant decoherence mechanisms, namely collisions with residual gas molecules and photon recoil heating, which ultimately set the timescale over which coherent expansion is possible.

With that in hand, Chapter 3 will move over to the setup used for the experiments, focusing on the parts of the apparatus that are essential to make long free evolutions possible. We will describe the optical tweezer and the laser, the back-scattered homodyne detection that provides the most sensitive readout of the axial motion, and the three orthogonal pairs of electrodes that are used both for feedback cooling and for the DC compensation itself. The feedback-cooling chain – including the phase-locked loops and the electronics used to apply both the feedback and the DC voltages simultaneously – will also be described, since the stability and noise properties of this chain directly affect how well the compensation can work.

Chapter 4 will then detail the compensation procedure itself: starting with 1D voltage scans that measurethe background force along a single axis, extending this to 3D by characterising and correctingthe cross-talk betweenthe electrode axes, and then showing howthe resulting compensation voltages are refined iteratively for increasing free evolution times. We will also describethe charge calibration procedure used to determinethe particle charge and howthe environmental charge state can be manipulated via a tungsten filament at low pressure. Beyondthe compensation procedure, Chapter 4 will present measurements that speak tothe coherence ofthe expansion: reheating measurements that showthe voltage noise ofthe DC supplies is negligible at typical compensation voltages, and a recompression measurement demonstrating thatthe state returns to its initial extent after a short free evolution. Finally, we will showthe main experimental results –the mean energy increase andthe growth ofthe position standard deviation as a function of free evolution time – both of which are consistent with coherent, force-free expansion up to several hundred microseconds.

We will then close in Chapter 5 with a discussion of the current limitations, some next potential setups, remaining gaps to the protocol of Section 1.2, along with some ideas for improving the experiment.

2. Theory

We will introduce the minimum of theoretical background needed to understand the experiments later on in Chapter 4. We start by explaining why we can trap nanoparticles with light fields at all, then we will highlight the important parameters of the particle dynamics in the trap as well as during free evolution, noting some pitfalls and challenges along the way. To understand these dynamics, we also touch upon the theory necessary to model the particle's motion.

2.1. Optical trapping

Light is able to transfer momentum, leading to radiation pressure. In 1970 A. Ashkin first used these forces stemming from the light field to push micrometre-sized particles in gases and liquids around, confining them in the radial direction but accelerating them along the propagation direction [34]. Only a few years later, it was shown that a single tightly focused laser beam could suffice to trap particles in all three directions [35].
The underlying principle of optical trapping can be understood by considering a photon's path through media of different refractive indices; however, let us first clarify the regime in which our particles reside. Regarding size, there are two main distinctions, relative to the wavelength of the light: the Rayleigh regime, where the particle size is much smaller than the wavelength of the trapping light, and the Mie regime, where the particle size is comparable to or larger than the wavelength. In this thesis, we will focus on the Rayleigh regime, as our particles have a diameter of 156 nm, which is significantly smaller than the wavelength of the trapping laser (1064 nm). For the particle, the light is therefore a homogeneous electromagnetic wave. In this regime, the particle can be treated as a point dipole with a polarisability α that depends on its material properties and size [36]. When the particle is placed in a large-scale inhomogeneous light field, such as that of a focused laser beam, it experiences two primary forces: the scattering force and the gradient force. The scattering force arises from the momentum transfer of photons being inelastically scattered by the particle. This force acts in the direction of the wave vector $\vec{k}$ and is proportional to the intensity of the light field [37, 38]. The gradient force, on the other hand, results from the interaction of the induced dipole moment of the particle with the spatial gradient of the light intensity. This force pulls the particle towards regions of higher intensity, drawing it towards the focus of the laser beam, where the intensity is maximal.
An illustrative sketch of the particle trapped in the optical field is shown in Figure 2.1, also indicating the axis labels x,y,z as well as the names "radial" and "axial" for the respective directions.

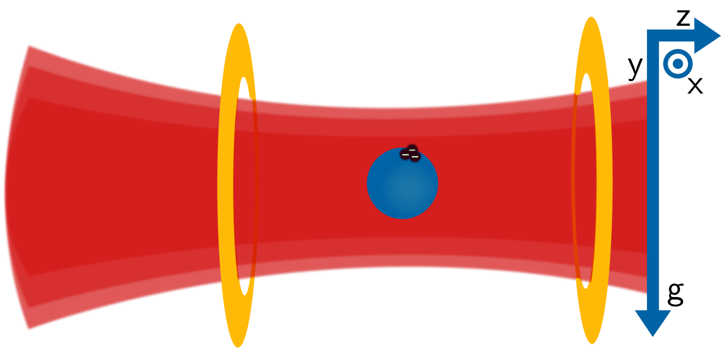

Figure 2.1.: **Schematic of the optical trapping setup.** We indicate the direction of gravity with an arrow labelled "g"; the laser propagates from left to right. The coordinate system is displayed to show the directions (x and y (radial), z (axial)) in our configuration. For clarity, although equipped with electrode pairs in all three directions, only the axial (z) electrode pairs are shown. The lenses and the electrode holder are also not shown.

Following [39] in our derivation, the total force F acting on a dipole (approximated as two masses m_1 and m_2 separated by a distance $\mathbf{d} = \mathbf{x_1} - \mathbf{x_2}$ resulting in a dipole moment $\mathbf{p} = q\mathbf{d}$) in the Rayleigh regime is

$$\mathbf{F} = (\mathbf{p} \cdot \nabla)\mathbf{E} + \dot{\mathbf{p}} \times \mathbf{B} + \dot{\mathbf{x}} \times (\mathbf{p} \cdot \nabla)\mathbf{B}, \tag{2.1}$$

where all bold quantities are vectorial and dependent on $\mathbf{x}$ and t. This is a general relation that holds for multiple regimes and incoming E and B fields. We can still drop terms by using the fact that the velocities we are considering are far slower than the speed of light c, and end up with a concise form for the force

$$\mathbf{F} = (\mathbf{p} \cdot \nabla)\mathbf{E} + \mathbf{p} \times (\nabla \times \mathbf{E}) + \frac{\mathrm{d}}{\mathrm{d}t}(\mathbf{p} \times \mathbf{B}) \tag{2.2}$$

$$= \sum_i p_i \nabla E_i + \frac{\mathrm{d}}{\mathrm{d}t}(\mathbf{p} \times \mathbf{B}), \tag{2.3}$$

for Cartesian indices $i = x, y, z$. As we are not interested in the instantaneous force, we can additionally take a time average denoted by $\langle \dots \rangle$ to which the last term does not contribute. As the final step, we look only at monochromatic fields and approximate to first order the dipole moment

$$p_i = \sum_i \alpha_{ij} E(\mathbf{x}), \tag{2.4}$$

with α_{ij} being the complex polarisabilities of the particle. For isotropic particles, $\alpha_{ij} = \alpha\delta_{ij}$.

The time-averaged force then turns out to be

$$\langle \mathbf{F} \rangle = \frac{\alpha'}{4} \nabla |\mathbf{E}|^2 + \sum_{i,j} \frac{\alpha''_{ij}}{2} \mathrm{Im}\left(E_i^* E_j\right), \tag{2.5}$$

with the first half being the gradient force and the second, non-simplified half being the scattering force.

The two radial directions are in principle rotationally symmetric, but if one takes into account the polarisation of the trapping light, this symmetry is broken and the two radial modes split in frequency [40]. An outline of the intensity profile calculated as in [39, Eq. 3.66, Ch. 3.6] follows and a sketched profile is shown in Figure 2.2.
For a linearly polarised beam (along x) focused by a high numerical aperture (NA) lens, the electric field in the focal region is given by:

$$\mathbf{E}(\rho, \phi, z) = -\frac{ikf}{2} \begin{pmatrix} I_0 + I_2 \cos(2\phi) \\ I_2 \sin(2\phi) \\ -2iI_1 \cos(\phi) \end{pmatrix} \tag{2.6}$$

where the integrals I_0, I_1, I_2 are defined as:

$$I_0(\rho, z) = \int_0^\alpha \sqrt{\cos\theta} \sin\theta (1 + \cos\theta) J_0(k\rho \sin\theta) e^{ikz\cos\theta} d\theta \tag{2.7}$$

$$I_1(\rho, z) = \int_0^\alpha \sqrt{\cos\theta} \sin^2\theta J_1(k\rho \sin\theta) e^{ikz\cos\theta} d\theta \tag{2.8}$$

$$I_2(\rho, z) = \int_0^\alpha \sqrt{\cos\theta} \sin\theta (1 - \cos\theta) J_2(k\rho \sin\theta) e^{ikz\cos\theta} d\theta. \tag{2.9}$$

Here J_n is the Bessel function of order n, α is the semi-aperture angle defined by the NA ($NA = n \sin\alpha$) and θ is the last remaining angular coordinate from the spherical coordinate system used to derive these relations. For the details of the derivation, we refer to [39] as this would be beyond the scope of this thesis.

Expanding the electric field around the equilibrium position of the trapped particle and ignoring terms above the second order [39, Eq. 15.73, Ch. 15.5], we find a potential

$$V(x, y, z) = -\frac{1}{4}\alpha' |\mathbf{E}(x, y, z)|^2 \approx \frac{1}{2} m \left(\Omega_x^2 x^2 + \Omega_y^2 y^2 + \Omega_z^2 z^2\right), \tag{2.10}$$

harmonic in all three directions, where Ω_i are the trap frequencies along the respective directions. In this harmonic approximation, all three motional degrees of freedom are decoupled and can be treated independently. This harmonic approximation is valid close to the focus, but for larger displacements, the potential becomes anharmonic, which is relevant for the "Big range z scan" method described in Section 4.3.1.
As discussed earlier, the trapping potential originates from the induced dipole moment within the particle. If we assume the polarisability α is uniform across all spatial axes, this directly leads to the intermediate formulation shown in Equation (2.10).

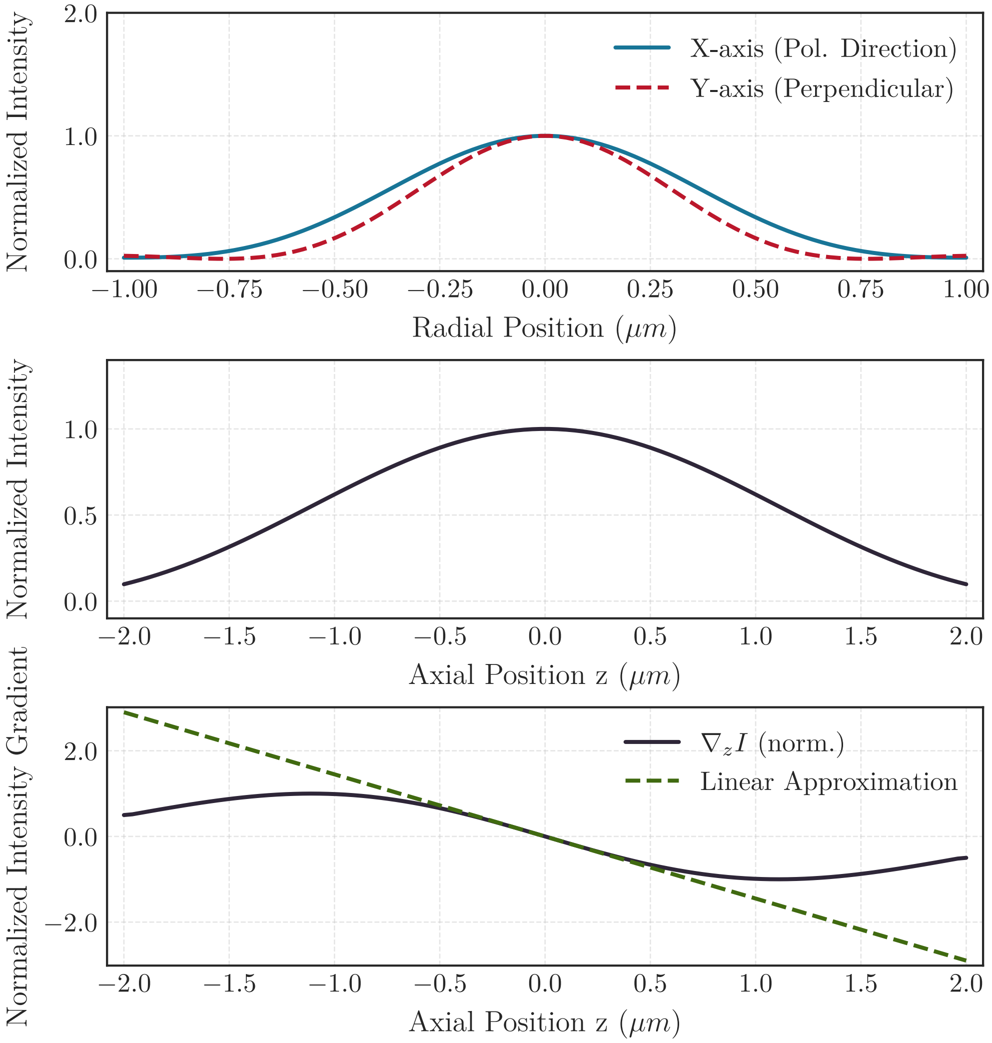

Figure 2.2.: **Intensity profile of a focused Gaussian beam, showing the radial and axial intensity distributions. Top:** Intensity distribution for the radial x/y axes at $z = 0$. **Middle:** Intensity distribution along the axial z-axis at $x = y = 0$. The particle is pulled towards the focus, where the intensity is highest. Due to radiation pressure, there is an additional slight offset along the propagation direction (z-axis) for the actual trapping position. **Bottom:** This shows the gradient of the intensity along the z-axis as well as the linear approximation around the focus.

Turning to measurement, the same process that gives rise to the scattering force also provides the signal used to infer the particle's position: each scattered photon carries momentum $p = \hbar k$ and thus conveys information about the particle's location. If N photons are collected, the shot-noise–limited position uncertainty decreases as

$$\Delta x \leq \frac{1}{\sqrt{2Nk}}. \tag{2.11}$$

Increasing the optical power increases N and therefore reduces this imprecision, but it also raises the rate of random momentum kicks from individual photon scattering events. These random kicks constitute radiation-pressure shot noise (RPSN), a back-action noise that perturbs the particle's motion. Hence there is a trade-off: higher power improves measurement precision but also increases back-action, which ultimately limits the net information that can be extracted about the particle's position [41]. Since in this thesis, we are not directly concerned with optimal cooling of the particle motion, nor with a certain occupancy we need to reach, we will not go deeper into this topic. To achieve the best possible detection, as a good practical rule, the optical power is kept as high as possible without saturating the detectors, so that the noise floor comes from the light shot noise and not from other electronic noise. Otherwise the upper limit for the cooling would be set by electronic noise and we would "waste" an unnecessary amount of particle position information, as the particle's motion can only ever be cooled to the total noise floor.

With these optical forces now in mind, we can understand the trajectory of the particle in the optical trap. The assumption of small displacements and consequently the harmonicity we found in Equation 2.10 allows us to treat the system as three non-interacting 1D harmonic oscillators, so we can focus on $z(t)$, with $x(t), y(t)$ behaving the same. The centre of mass (COM) motion of the particle in the z direction, $z(t)$, can be described by the Langevin equation for a damped harmonic oscillator with a deterministic (F_{Det}) and fluctuating force term $F_{\text{Fluct}}(t)$:

$$\ddot{z}(t) + \gamma \dot{z}(t) + \Omega_z^2 z(t) = \frac{1}{m}\left(F_{\text{Det}} + F_{\text{Fluct}}\right) \tag{2.12}$$

As is done in [41], the damping rate γ can be further decomposed into constituents

$$\gamma = \gamma_{\text{thermal}} + \gamma_{\text{radiation}} + \gamma_{\text{feedback}} \tag{2.13}$$

where γ_{thermal} accounts for the particle's interaction with the background gas, $\gamma_{\text{radiation}}$ describes the interaction with the light field, and γ_{feedback} represents the damping introduced by cold damping feedback cooling. A more thorough discussion of these terms and their relative importance in different regimes can be found in Section 2.5.
A subtle and important point to note here is that in this thesis we need to make a distinction between *trapping* and *levitating* that most of the literature does not (need to) make. From now on, we will use the term *trapping* to refer to the situation where the particle is confined by the harmonic potential of the optical trap, and *levitating* to refer to the situation where the particle is freely evolving without the influence of the trap,

but suspended against gravity by the electrostatic forces we intentionally apply. From Earnshaw's theorem [42], we know that stable trapping of a charged particle in a static electric field is not possible, but levitation is possible for limited times.

2.2. Squeezing

During free evolution, the trap is turned off, meaning that the particle is no longer confined by the harmonic potential. In this case, assuming no energy is added through decoherence, the particle's position and momentum will evolve according to the free particle Hamiltonian $H_0 = \frac{\hat{p}^2}{2m}$, leading to a coherent increase in the position variance over time.
As the particle's COM motion is in a thermal state, it is initially represented by a Gaussian distribution in phase space. There are two parameters that describe this distribution: the first moment, which is the mean position and momentum, and the second moment, which is the covariance matrix describing the shape and size of the distribution in phase space. The first moment after a free evolution time τ can be expressed as

$$\begin{pmatrix} z(\tau) \\ p(\tau) \end{pmatrix} = S(\tau) \begin{pmatrix} z(0) \\ p(0) \end{pmatrix} = \begin{pmatrix} 1 & \tau \\ 0 & 1 \end{pmatrix} \begin{pmatrix} z(0) \\ p(0) \end{pmatrix}. \tag{2.14}$$

This first moment captures the trajectory of the centre of the phase space distribution, making it equal to the classical trajectory of a free particle.

The second moment, which describes the shape and size of the distribution in phase space, evolves according to the symplectic[1] transformation

$$\mathbf{\Sigma}(\tau) = S(\tau)\mathbf{\Sigma}(0)S^{\mathrm{T}}(\tau), \tag{2.15}$$

where $S(\tau)$ is the evolution matrix defined in Equation (2.14). For a covariance matrix $\mathbf{\Sigma}$ composed of elements $\Sigma_{ij} = \mathrm{cov}(\hat{i}, \hat{j}) = \frac{1}{2}\langle \hat{i}\hat{j} + \hat{j}\hat{i}\rangle - \langle \hat{i}\rangle\langle \hat{j}\rangle$, this operation yields

$$\begin{pmatrix} \Sigma_{zz}(\tau) & \Sigma_{zp}(\tau) \\ \Sigma_{zp}(\tau) & \Sigma_{pp}(\tau) \end{pmatrix} = \begin{pmatrix} \Sigma_{zz}(0) + 2\tau\Sigma_{zp}(0) + \tau^2\Sigma_{pp}(0) & \Sigma_{zp}(0) + \tau\Sigma_{pp}(0) \\ \Sigma_{zp}(0) + \tau\Sigma_{pp}(0) & \Sigma_{pp}(0) \end{pmatrix}. \tag{2.16}$$

We have used a Weyl-symmetrized definition of the covariance, such that the covariance matrix is always symmetric and real, and then calculated Equation (2.15) with the definition of $S(\tau)$ from Equation (2.14).

The phenomenon of squeezing comes from the growth of the off-diagonal terms. As the free expansion time τ increases, the covariance $\Sigma_{zp}(\tau)$ also increases, leading to a growing correlation between position and momentum, since components of the state with higher

[1] Symplectic for phase space is like unitary for states, meaning in our specific case we need it to be symplectic to make sure Liouville's theorem holds and that it preserves symmetries of the $\hat{x}$ and $\hat{p}$ commutation relation.

momentum will travel proportionally further in space. In phase space, this corresponds to a shearing effect. If we start with a state of an oscillator with frequency Ω, represented by a circle, the free evolution of time τ will tilt and stretch this distribution along a diagonal axis with angle $\theta(\tau)$ to the position axis, which can be calculated as

$$\theta(\tau) = \frac{1}{2}\arctan\left(\frac{2}{\tau\Omega}\right). \tag{2.17}$$

From Liouville's theorem, we know that the total energy of this phase space ellipse must be conserved, that is to say that any increase in the area is related to a transfer of energy. The volume of such a phase space ellipsoid is proportional to $\sqrt{\det(\mathbf{\Sigma})}$, sometimes called a *generalised variance* [43, p. 264], so as the distribution undergoes this extreme stretching along its major principal axis, it is geometrically necessary to narrow along its orthogonal minor axis. This compression is called *squeezing*, and is especially interesting once the squeezing becomes so significant that the squeezed quadrature (ellipse axis) is smaller than the initial ground state variance.

The symplectic transformation $\mathbf{\Sigma}(\tau) = S(\tau)\mathbf{\Sigma}(0)S^{\mathrm{T}}(\tau)$ we have discussed above describes an integrated, single macroscopic jump of the covariance matrix from time 0 to τ for a perfectly isolated quantum system. To correctly integrate rates of energy exchange, this simpler transformation is not sufficient. Consequently, one has to use the Lyapunov equation, a continuous-time, differential formulation of this exact same evolution, but due to its differential nature, heating and cooling effects can be included. While the rigorous treatment is beyond the scope of this thesis, for a system with drift matrix $\mathbf{A}$ (encoding the deterministic dynamics) and diffusion matrix $\mathbf{D}$ (encoding stochastic heating), the covariance evolves as [44, 45]

$$\dot{\mathbf{\Sigma}} = \mathbf{A}\mathbf{\Sigma} + \mathbf{\Sigma}\mathbf{A}^{\mathrm{T}} + \mathbf{D}. \tag{2.18}$$

For example, during free evolution with gas damping γ and momentum diffusion coefficient D_{pp},

$$\mathbf{A}_{\mathrm{free}} = \begin{pmatrix} 0 & 1/m \\ 0 & -\gamma \end{pmatrix}, \qquad \mathbf{D}_{\mathrm{free}} = \begin{pmatrix} 0 & 0 \\ 0 & D_{pp} \end{pmatrix}. \tag{2.19}$$

In Section A.3 we will discuss in more detail the model used to add theory comparison to the measurements according to the Lyapunov equation.

2.3. Trap (Duffing) nonlinearity

What is important to keep in mind and part of a measurement that will be discussed in Section 4.3.1 is that the particle is not actually trapped in a harmonic potential. As mentioned in Section 2.1, the light in reality creates a Gaussian potential, which is approximately harmonic close to the centre, but becomes more and more anharmonic for larger displacements. Since it is not within the scope of this thesis to derive the exact form of the potential, we will just give a sketch of the form of the potential and its consequences for the particle's dynamics, following [46, ch. 2]. The approximation of the potential now gains terms additional to those in Equation (2.10), leading to a potential of the form (given in [46, equation 2.34] [2]):

$$\begin{aligned} U_{\text{opt}}(\mathbf{r}) \approx -U_0 + \frac{m}{2}\left(\Omega_x^2 x^2 + \Omega_y^2 y^2 + \Omega_z^2 z^2\right) \\ -\frac{m^2}{8U_0}\left(\Omega_x^4 x^4 + \Omega_y^4 y^4 + 2\Omega_z^4 z^4\right) \\ -\frac{m^2}{4U_0}\left(\Omega_x^2\Omega_y^2\, x^2y^2 + 2\Omega_x^2\Omega_z^2\, x^2z^2 + 2\Omega_y^2\Omega_z^2\, y^2z^2\right). \end{aligned} \tag{2.20}$$

In Equation (2.20) we see two new types of terms and the newly introduced parameter U_0 for the trap depth: the quartic terms in the second line, which are nonlinear additions in each axis but still independent. We also see cross-talk terms in the third line, which couple the different axes. These *cross-coupling* terms are the reason our independent oscillator model breaks down here, making this regime further undesirable for experiments, since this then means that modelling and readout of the particle's state become increasingly complex. We also note that the trap depth plays a more important role, as for deeper (stronger) traps, the nonlinearities show up later, meaning that (for the same given displacement) in that case we expect them to be weaker.
Neglecting the cross-coupling terms for simplicity, this means that the particle will now have an *amplitude-dependent* trap frequency

$$\Omega_{\text{z,NL}} = \Omega_{z,0} + \Delta\Omega_{\text{NL}} = \Omega_{z,0} + \frac{3}{8}\Omega\xi A^2. \tag{2.21}$$

which decreases with increasing amplitude of oscillation A, given a $\xi < 0$ [46], which is then aptly called a weakening nonlinearity [3]. Physically, this can be understood by imagining the particle going around in the actual Gaussian potential compared to the harmonic approximation. As it enters the nonlinear part, the potential becomes shallower, so the system will explore more space. These larger oscillations will take more time, reducing the effective trap frequency.

2.4. Free evolution and its parameters

Building on the necessary background of optical trapping and particle dynamics, we can continue toward the more applied and experimental aspects by introducing the two main figures of merit that we will use to evaluate the success of our free evolution experiments.

[2] in terms of trap depth or trap waist in the respective directions ω_{0i} but rewritten here in terms of the frequencies Ω_i. This can be converted by using $\Omega_i = \frac{2}{w_{0i}}\sqrt{\frac{U_0}{m}}$ and comparing like terms of [39, 46]

[3] Maybe worth mentioning is that $\xi > 0$, so a stiffening nonlinearity, is also physically possible and present in other systems.

2.4.1. Figure of Merit 1: Position Variance

As the overall goal of this thesis is to find a method to allow for state size expansion, we will use the position variance as a final figure of merit for this. Releasing a non-squeezed thermal state with initial mean occupancy $\bar{n}_z$ for a time τ will coherently increase the axial position variance according to

$$\sigma_z^2(\tau) = z_{\mathrm{zp}}^2(2\bar{n}_z + 1)(1 + \Omega_z^2\tau^2), \tag{2.22}$$

not including any decoherence or heating effects. Here, z_{zp} is the zero-point position uncertainty of the particle in the trap, and Ω_z is the trap frequency along the z-axis. This equation shows that the position variance increases quadratically with time during free evolution, being a key aspect of how to achieve large state expansion. The initial position variance is determined by the initial mean occupancy $\bar{n}_z$ and the zero-point uncertainty, and the term $(1 + \Omega_z^2\tau^2)$ captures the coherent increase in variance due to free evolution.
Although this figure of merit is the main focus and motivation, we usually are unconcerned about the position variance at first, as it leaves no direct pathway for improvement of the measurement. Of course, one can estimate the evolution of the state variance by taking the variance of the fitted trajectories after recapture, but during the experiment, we usually instead look at a different figure of merit that is more directly linked to the external forces acting on the particle during free evolution: the mean energy increase.

2.4.2. Figure of Merit 2: Mean Energy Increase

In addition to the position variance, we will also use the mean energy (or rather the increase) as a measure of how well we mitigate external forces – the main issue for free evolution of charged nanoparticles.
They are therefore closely related, as the lowest achievable mean energy increase is that of a perfectly force-free evolution, leading to – theoretically – indefinite evolution times and with that arbitrarily large state expansion.
From [30], we use the following equation for the mean energy E_z increase after a free expansion time τ in the presence of a static force F_0.

$$\langle E_z \rangle = E_z(0)\left[1 + \frac{\Omega^2\tau^2}{2}\right] + F_z^2\frac{\tau^2}{2m}\left[1 + \frac{\Omega^2\tau^2}{4}\right], \tag{2.23}$$

where Ω is the trap frequency along z and m is the particle mass. $E_z(0)$ is the initial mean energy before the free expansion. As the modes can be treated independently, this equation also holds for x and y directions and provides the fundamental scaling laws that will be used in this thesis: The F_{i}^2 scaling is observed during the compensation scans later on, while the scaling with τ is used in residual force measurements after optimal compensation. For force-free evolution, the mean energy should scale quadratically with τ.

To better understand what the individual contributions to this mean energy increase are, we will rewrite Equation (2.23) as

$$\langle E_z \rangle = \underbrace{E_z(0)}_{\text{⚀}} + \underbrace{\frac{E_z(0)\Omega^2}{2}\tau^2}_{\text{⚁}} + \underbrace{\frac{F_z^2}{2m}\tau^2}_{\text{⚂}} + \underbrace{\frac{F_z^2\Omega^2}{8m}\tau^4}_{\text{⚃}} \tag{2.24}$$

and go through the terms one by one:

- ⚀ **Initial energy**: This term represents the initial mean energy of the axial COM motion before the free evolution begins. It is determined by the thermal state of the particle in the optical trap, which is influenced by factors such as feedback cooling and environmental conditions as further treated in Section 2.5. Since during the free evolution nothing changes the initial energy, this term remains constant throughout the process, by definition.

- ⚁ **Potential energy contribution from initial velocity**: This term can be interpreted as an increase in the mean energy due to the particle's initial velocity (proportional to $\sqrt{E_z(0)}$) at the moment the trap is turned off. Without any external force, the particle will continue to move with this initial velocity during the free evolution time τ. As a result, its position variance increases, leading to a corresponding increase in potential energy when the trap is re-engaged, due to displacement in $V(z)$.

- ⚂ **Kinetic energy contribution from the force**: Since a constant force F_z acts on the particle during the free evolution, it imparts an acceleration to the particle during that time. This acceleration leads to an increase in the particle's momentum, and consequently, its kinetic energy increases.

- ⚃ **Potential energy from the force**: This term represents the increase in potential energy due to the displacement of the particle caused by the constant force F_z during the free evolution time τ. As the particle is accelerated by the force, it moves away from its initial position, leading to a change in its potential energy when the trap is re-engaged.

2.5. Temperature, Heating and Decoherence

The particle has a certain steady-state energy E_∞ and, given that the state is thermal, can be viewed from a thermodynamic perspective as having an effective COM temperature T according to the equipartition theorem: $E_\infty = k_\mathrm{B}T$, with k_B being Boltzmann's constant. The particle initially gets captured in the optical trap, surrounded by background gas and is therefore strongly coupled to its thermal environment, at temperature T. As the chamber is evacuated, the coupling to the background gas decreases. Nevertheless, we continue to use the effective temperature to characterise the motion. The internal degrees of freedom and their temperature are ignored in this view.
To transition to a quantum mechanical description of the particle's state, it is convenient to use the mean phonon occupancy n. For a quantum harmonic oscillator with frequency Ω_0, the energy is quantised in units of $E = \hbar\Omega_0$, where $\hbar$ is the reduced Planck constant. Therefore, we can express a relation between the energy in the high-temperature limit and the quanta of energy, the mean occupancy n_th, as

$$n_\mathrm{th} = \frac{k_\mathrm{B}T}{\hbar\Omega_0}. \tag{2.25}$$

These two descriptions of the energy are not yet reconciled with each other. We can remedy this by rewriting the steady state energy as the sum of the Bose-Einstein statistics and the zero-point energy, leading to

$$E_\infty(\Omega_0, T) = \frac{\hbar\Omega_0}{\exp(\hbar\Omega_0/k_\mathrm{B}T) - 1} + \frac{\hbar\Omega_0}{2}, \tag{2.26}$$

where the second term is the zero-point energy [39, ch. 15].
We see that for $k_\mathrm{B}T \ll \hbar\Omega_0$ this expression is of the form $E_\infty|_Q = \hbar\Omega_0/2$, so it reduces to the zero-point energy. In the opposite limit, where $k_\mathrm{B}T \gg \hbar\Omega_0$, the energy is once again the thermal energy $E_\infty|_T = k_\mathrm{B}T$ [39, ch. 16.1].

In addition to the damping forces discussed in Equation (2.12), various processes *add* energy to the system, leading to a change in the steady state occupancy n_∞. Following [41] and Equation (2.13), we find

$$n_\infty = \frac{\Gamma_\mathrm{thermal} + \Gamma_\mathrm{recoil} + \Gamma_\mathrm{feedback}}{\gamma_\mathrm{thermal} + \gamma_\mathrm{radiation} + \gamma_\mathrm{feedback}}, \tag{2.27}$$

i.e., a ratio of heating rates Γ and damping rates γ. The heating rates come from noise in both the feedback loop and DC-bias voltages ($\Gamma_\mathrm{feedback} = \Gamma_\mathrm{cooling} + \Gamma_\mathrm{bias}$), recoil from scattered photons (Γ_recoil) and collisions with the gas (Γ_thermal). In our experiments, when the trap is on, the dominant source of heating comes from the recoil of scattered photons.

This steady state n_∞ describes the equilibrium occupation number if the particle is actively cooled and therefore represents the state just before it is released for a free evolution. However, during this free evolution, the feedback is turned off and the trapping light is off, so Γ_{feedback} and Γ_{recoil} are negligible (given proper consideration to the choice of electronics) and zero respectively, and the only remaining source of decoherence is from collisions with the background gas. This means that the particle's state will decohere due to these collisions, and it will lead to an increase in the mean energy and position variance over time. The rate of this decoherence can be characterised by Γ_{thermal}, which depends on the pressure of the background gas and the size of the particle. At gas pressures below 10 mbar, the damping γ to the nanoparticle is linear in gas pressure [41]. This rate is also vital for future quantum experiments, as a collision with a gas molecule will locate the particle and thereby limit the timescales for experiment [47].

The only remaining electronic noise source during the free evolution time is the DC bias which is also included in Γ_{feedback}. While not negligible by default, in Section 4.1 we measure at what voltage the electronics used add significant noise to the system, verifying that they do not add significant noise at relevant levels. Therefore, in a free evolution, with the light off, we are in a regime where Γ_{recoil} and Γ_{feedback} are negligible, and Γ_{thermal} is the dominant source of decoherence, meaning the particle is essentially only decohered by collisions with the background gas. We will further measure this decoherence in Section 4.2.

3. Experimental Setup

In this chapter, we will describe the experimental setup used and developed throughout this thesis' work. Even though the used setup and components are standard, it will still benefit the understanding of the components and challenges discussed later. We will first give a lay of the land in Section 3.1, describing the main components of the setup and their purpose. Then, in Section 3.2, we will cover the control equipment. The optical setup, vacuum chamber and bulk of the electronics were already in place at the start of this work and built by Y. Y. Fein. The model/parts numbers are only listed in Section A.6 in the appendix and only the necessary properties of the equipment are described in the main text.

3.1. Setup Overview

For the sake of readability, we will divide the setup will be divided into multiple parts that serve different purposes. The subsequent chapter titles correspond to the colour-coded parts in Figure 3.1 but it is important to acknowledge the somewhat arbitrary choice of "region" boundaries.

3.1.1. Laser

The laser used in this experiment is a Coherent Mephisto operating at a wavelength of 1064 nm. An acousto-optic modulator (AOM) in double-pass configuration allows for fast switching of the laser power. The AOM is driven by a radio-frequency (RF) signal from a Wieserlabs RF drive, whose signal is sent to a Mini-Circuits amplifier, a Mini-Circuits switch, and then to the AOM. The TTL pulse for switching the Mini-Circuits switch comes from a Keysight arbitrary waveform generator, which can thus control the timing of the laser pulses. After the AOM, the beam passes preparation optics, and a half-wave plate, then a polarising beam splitter (PBS) to split off a part of the beam for the local oscillator (LO, see Section 3.1.3) of the homodyne detection. The other part of the beam is then the trapping beam.

3.1.2. Trap

After the beam preparation and LO pick-off, the trapping beam passes through a Faraday rotator (FR) and two waveplates (a half-wave (HWP) and a quarter-wave plate (QWP)) before entering the vacuum chamber. This allows us to adjust the polarisation of the beam, which is important for setting the trap polarisation and thereby the detuning of the radial modes with respect to each other as well as to spin up librations and rotations [48]. This is explained in more detail in Section 2.1. The FR is used to separate the back-scattered (BW) light from the main trapping beam for the homodyne detection, as described in Section 3.1.5.

The main part of the setup is the optical trap itself, which is located inside a vacuum chamber. The laser beam is focused by an aspheric lens (Asphericon) with a numerical aperture (NA) of 0.8 into the centre of the vacuum chamber, creating a strong optical potential that can trap the nanoparticles.

As discussed in Section 2.1, the trapping potential is created by the gradient force of the focused laser beam. The high NA lens is crucial for achieving a sufficiently strong gradient force to overcome other forces acting on the particle, such as radiation pressure and gravity. After the optics before the vacuum chamber split off parts of the light and change the power, in standard operation, 600 mW of power is entering the vacuum chamber and formms the trapping potential.

To get particles into the trapping region in the first place, a nebuliser diffuses a solution of silica nanoparticles in isopropanol into a fine mist. This mist is then guided into the vacuum chamber through a thin tube. A particle gets trapped in the focus of the beam at random at slightly below atmospheric pressure. Reloading typically starts at low pressures following particle loss, meaning the air that passed through the nebuliser constitutes the primary source of gas and aerosols in the vacuum chamber, saturating the environment with particles. This results in a high probability of trapping a particle, but it also increases the likelihood of trapping clusters or isopropanol droplets instead of single nanoparticles.

This is not ideal, but it is a consequence of the trapping method and can be detected by looking at how much laser light the particle scatters. For this, we can either route the BW light to a power meter or look at the image of the particle that is taken by a camera running a live monitoring feed. The camera looks through the chamber from the side, along the x-axis. After trapping, the pressure is lowered, which then evaporates the isopropanol, leaving a single silica nanoparticle in the trap.

3.1.3. Local Oscillator

To create the Local Oscillator for the homodyne detection, a part of the main laser beam is picked off before it enters the vacuum chamber. This ensures that the LO and the signal beam are phase-coherent, which is a requirement for homodyne detection. The LO beam then passes through a set of waveplates (a half-wave and a quarter-wave plate), is coupled into a multimode fibre, runs through a fibre beamsplitter (BS) being combined with the backscattered signal beam from the trap and is then sent to the balanced homodyne detector as explained in Section 3.1.5. To make sure the phase between the LO and signal beam is optimal, the path length of the LO beam can be adjusted with an electro-optical modulator (EOM) that locks the phase of the two arms, by taking in the monitor outputs of the balanced photodiode and feeding it into a PID board that drives the EOM.

3.1.4. Electrodes

The lenses that focus the beam into the vacuum chamber are mounted on a machined electrode assembly, which also holds the ring electrode pairs for applying electric fields to the particle. The CAD model of this assembly is shown in Figure 3.2. The electrodes are used for compensating stray electric fields and for applying feedback forces to the charged particle for cooling. The electrodes are electrically insulated by being mounted in a piece of PEEK, a vacuum-compatible insulator. The direction of the axes is also indicated in Figure 3.2, matching with the earlier sketch of the trapping laser (Figure 2.1). Due to the geometry of the setup, the electrodes for the x and y directions have a roughly $10\times$ worse force transduction compared to the z direction, which is important to keep in mind for the later methods and results. This comes from both the distance, as well as the different radius of these electrode pairs. The electrodes are connected to the control electronics outside of the vacuum chamber via a feedthrough, which allows for a vacuum-compatible connection to the inside.

3.1.5. Detections

For the actual data acquisition and cooling, multiple methods are employed. First, a camera is used to visually monitor the trap, cables and the particle. This looks into the chamber from the side, which would correspond to the x-axis in Figure 3.2. This is not used for any measurements, but is very helpful for aligning the trap and making sure the particle is still trapped, as well as for estimating the particle size and makeup by analysing the scattered light pattern, as sometimes clusters or droplets are trapped instead of single nanoparticles. A picture of the forward detection is included in Figure A.7

Primary (z) Axis Detection

A photodiode is used for z-detection, which has the forward-directed light centred on it. This photodiode is easy to align, has a sufficient range and is mounted in the detection box and passes through a HWP + PBS combination to split off part of the light for the radial detections. While the forward detection is not as good as a backward homodyne detection, its robustness makes it the detection of choice for pump-down and as a backup detection while changing the homodyne detection at low pressures. Additionally, it is used for leaving the setup running overnight, as it is not prone to phase-lock range issues that can occur with the homodyne detection due to the locking PID integration term reaching its limits.

The homodyne detection is comprised of the light that is backscattered from the particle, which is separated from the main trapping beam by the Faraday rotator (FR) after it exits the vacuum chamber again through the same port upon reflection off the particle. This backscattered light is combined with the LO beam on a fibre beamsplitter, and the two output ports are directed onto a balanced photodiode. As the polarisation of the LO and signal beam must also be matched, a set of waveplates (a half-wave and a quarter-wave plate) is placed in the LO and signal arms before the two beams are combined. Since the balanced photodiode measures the difference between the two photodiodes and then amplifies the result, it is very sensitive to small changes in the interference between the LO and signal beam, allowing for high-sensitivity detection of the particle's position along the optical axis (z-axis).

It can be shown that there is a disparity in information content between the backscattered and forward-scattered light, making backscattered light (BW) fundamentally better suited for sensitive detection [49, 50]. Furthermore, due to how the detectors work, the available signal for cooling has a lower noise floor compared to the forward-scattered (FW) detection. Therefore, BW is primarily used for feedback cooling of the z mode at low pressures.

The signals from the two schemes are then processed further for cooling and data acquisition, as described in Section 3.1.6.

Radial (x and y) Axis Detection

Light that is forward-scattered (FW) and not used for the FW z axis detection is split up by two further WP and PBS combinations to obtain two independent radial detections composed of the same original light. Power control of these is done by the WP and PBS for each detection. The particle signal is collected by using a split detection in each of these arms: The light is shone upon a D-shaped mirror, which splits the beam into two halves. Each half is then focused onto a fibre coupler, getting passed through a multimode fibre to the balanced detectors. The coupling into the fibres is used to obtain similar powers on both inputs of the balanced detector, which is crucial for direct use of the RF-Output of the detector, since the subtraction of the two signals requires them to be approximately balanced. Otherwise, the balanced detector will not output a useful signal and the cooling will not work. This can be partially overcome by using the digital bandpass filter and a PID controller as described in Section 3.1.6, but also there it is still important to have light coupled into the fibres such that the final processed signal enhances the wanted mode.

As the balancing drifts over time due to temperature changes and other factors, the DC outputs of the balanced detector arms are monitored and the optical components (the D-shaped mirror for x, the PBS for y) are adjusted manually to keep the DC level difference small enough to be in range for the detector to work. One fix for this was implemented after the data for this thesis was taken and therefore only affects how the detection works for the long free evolution (Section 4.11) and the squeezing measurements (Section 4.12). The fix was to bypass the balancing part of the detector altogether and use the monitor arms of the balanced detector to feed into two digital PIDs running on a FPGA with positive and negative proportional coefficients, respectively, which results in a feedback loop that is significantly more robust to drifts of the respective DC levels in the two arms of each detection. This represents a significant improvement as it uses the same hardware that is already needed for the more robust bandpass filter cooling as described in the next chapter.

In any case, however, the light that is coupled into the multimode fibres to then go to the detectors must be chosen such that the mode of choice is amplified and the other mode is suppressed by the subtraction. This works, since one of the two halves of the beam gets its phase shifted with respect to the other upon reflection off the D-shaped mirror. In the end, using pinholes and careful coupling, it can be achieved that the wanted detection has a peak of 50 dB, which can then be used to generate the cooling signal, while the other radial mode is suppressed to a level of a few tens of dB. As the axial motion is radially symmetric, the D-shaped mirror will split the beam into two halves that have the same information, but with a phase shift of π between them. This means that the subtraction of the two halves of the beam results in a signal that is amplified by a factor of 2, while the noise that is common to both halves (like laser intensity noise) is suppressed. While it would be preferable to only have one mode in each detection, this z mode cannot be eliminated and shows up in both radial detections. Practically, due to the difference in frequency $\approx 200\,\mathrm{kHz}$, the z mode is not a problem in terms of isolation.

3.1.6. Cooling

For the experiments presented in this thesis, the nanoparticle's centre-of-mass (COM) motion must be prepared in a cooled state. If left in thermal equilibrium with the surrounding gas, the particle undergoes continuous Brownian motion. These random collisions with gas molecules result in a thermally excited state, characterised by high mean motional energy. The primary goal of feedback cooling is to actively remove this energy, thereby reducing the particle's COM temperature for subsequent experiments.
Active cooling becomes particularly crucial at the low pressures required for high-precision measurements. While reducing the pressure is effective at isolating the nanoparticle from the thermal environment, it also removes the natural damping provided by gas collisions. Without this damping, the particle's motion is still susceptible to technical noise, which will cause the energy to increase, also leading to large oscillations and eventual loss from the trap as the energy grows greater than the potential depth. Feedback cooling introduces an artificial, controllable damping force that both stabilises the particle and actively reduces its motional energy.

This technique is referred to as cold damping because it circumvents the constraints of a conventional thermal bath. In a normal system, the damping force (dissipation) from a hot reservoir (like gas) is intrinsically linked to a random fluctuating force (heat), as described by the Fluctuation–Dissipation Theorem. In cold damping, this link is broken. The feedback force is engineered to be proportional to the particle's velocity (practically, a suitably phase-shifted position signal for a sinusoidal motion). This force adds significant damping without an associated thermal noise term. Consequently, the effective mechanical damping is increased, coupling the oscillator to a cold reservoir. This enables the cooling of the COM to temperatures far below what passive gas damping could achieve, limited in theory only by the noise in the measurement and feedback electronics and quantum mechanics [9, 49].

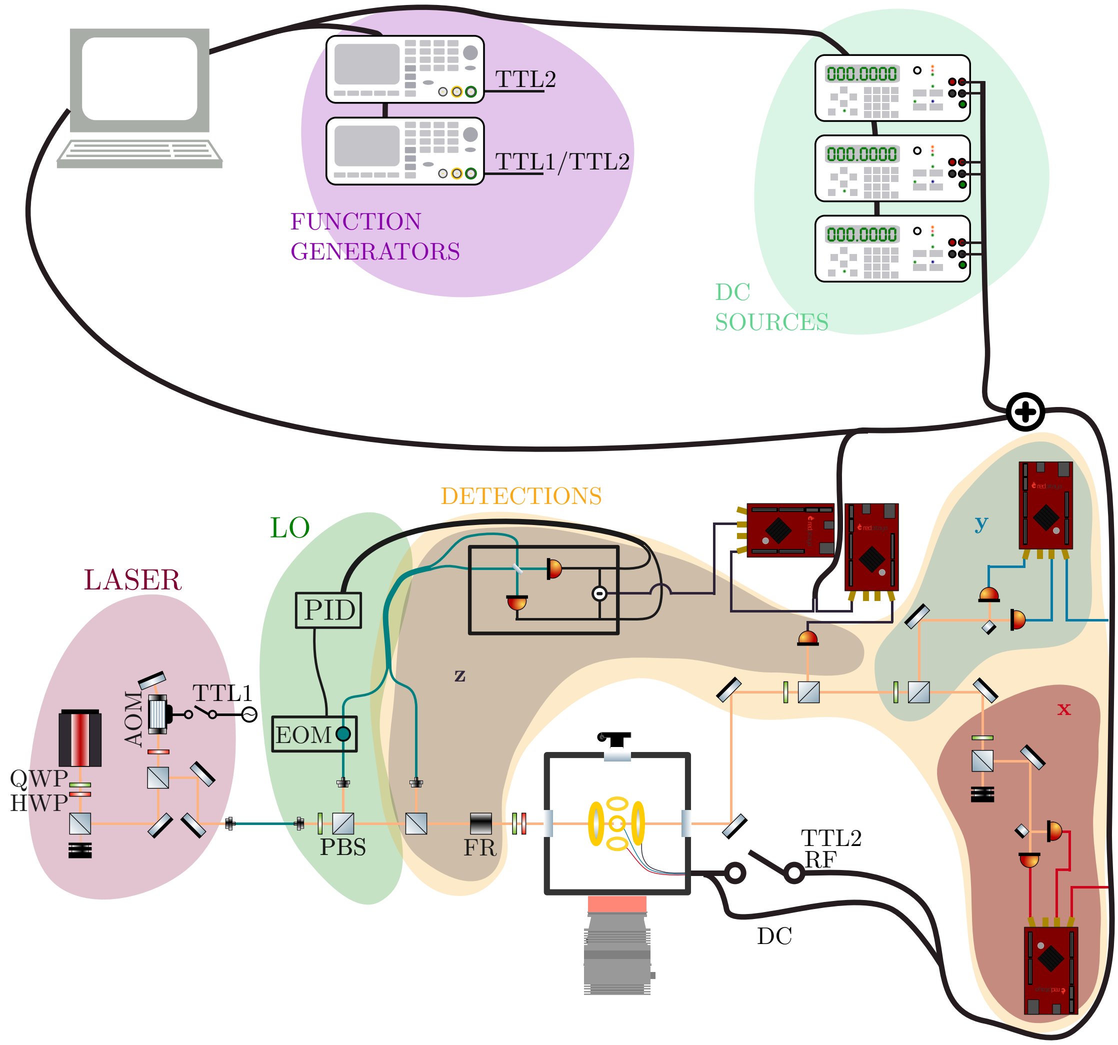

Figure 3.1.: **Schematic of the full experimental setup.** The main components of the setup are colour-coded and correspond to the subsequent chapter titles. The laser components from Section 3.1.1 are shown in red to the leftmost part of the image. Continuing to the right, the LO Pick off from Section 3.1.3 is indicated in green. In orange follow all the detection components described in Section 3.1.5, which are split into the forward (FW) () and backward (BW) z detections in deep purple, and the x and y radial detections in red and blue, respectively, following the convention of color mapping the axes in the other figures. The trap components of Section 3.1.2, vacuum pump and chamber enclosing the elecrodes from Section 3.1.4 are indicated by having no colour below all the other components. On the top, the control equipment described in Section 3.2 and Section 3.2 is indicated in lavender for the function generators and in mint green for the voltage sources. Cables and connections are shown partly, but they should be understood to be mulipurpose and indicate connections rather than the actual cables used. TTL connections are not connected, since they depend on the measurment protocol.

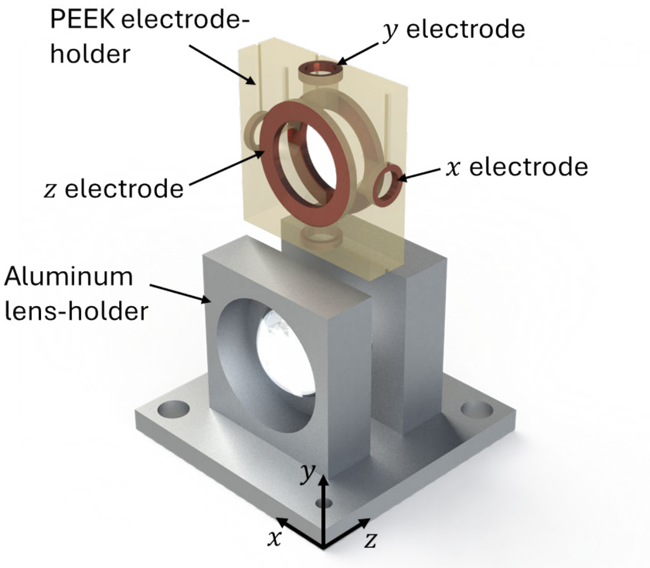

Figure 3.2.: **Electrode assembly CAD model**: Seen here is the electrode assembly that holds the lenses and the ring electrode pairs for the x/y/z directions. The electrodes are mounted in a piece of PEEK for electrical insulation. The direction of the axes is also indicated, to reference with the earlier sketch of the trapping laser (Figure 2.1). It can be seen that the electrodes for the x and y are further away from the trapping position and have a smaller radius, leading to a roughly 10× worse force transduction compared to the z direction. The electrodes are connected to the control electronics outside of the vacuum chamber via a feedthrough, which allows for a vacuum-compatible connection to the inside. (Image Credit: Y. Y. Fein)

In this work, the feedback loop is realised by measuring the particle's position in real-time with the already discussed detections and applying a corrective force via the electrode pairs. The high-speed signal processing is handled by Field-Programmable Gate Arrays (FPGAs). For the measurements, different cooling signal generation methods were used:

PLL-Cooling

A phase-locked loop (PLL) is a control system that generates an output signal with a minimised phase error between input and internal reference signal. It does this by continuously adjusting the frequency of an internal reference oscillator until the error is minimised [1]. In the locked state, the PLL has a phase error Δ_θ that is constant and a frequency error $\dot{\Delta}_\theta$ of zero. Given a sinusoidal input signal with varying frequency, the PLL can generate a clean output signal that is locked in phase to the input. For small deviations in the frequency [2] any mismatch in frequency also manifests in a phase difference and therefore the PLL locks both in frequency and phase. This is a crucial feature we will exploit in Section 4.4 and is the reason we can use phase-locked loops for the cooling of the particle.

As it is done digitally, the PLL can also additionally shift the phase-locked output with respect to the locked internal reference signal, intentionally phase-shifting the output to create a feedback signal that is proportional to the velocity (a $\pi/2$ phase shift for a sinusoidal signal) of the particle.

As it generates a sine wave from the instantaneous frequency of the particle signal (we can specify the cutoff frequencies of the phase-locked loop, so that it only tracks the signal at the desired frequency), the PLL should give one of the cleanest signals to use for cooling, as any other signal artefacts and noise at other frequencies are filtered out by the PLL, acting as a bandpass filter. However, in practice, we found that this is also one of the issues, since the PLL can lose lock on the particle signal, especially upon recapturing it after free evolutions where nonlinearities in the trap are explored (see Section 2.3), shifting the frequencies outside the specified bandwidth of the PLL. The PLL then generates a signal that is not related to the particle motion, which can lead to heating instead of cooling or just not cooling at all. This is especially problematic for the long free evolutions, where the particle signal could have explored nonlinear regions already and therefore changed its frequency, making it harder for the PLL to relock. This is one of the reasons we switched to the digital bandpass filter cooling method described in Section 3.1.6, which is more robust to changes in the particle signal frequency and, since digital, can be tuned live while cooling, just like the PLL. As the PLL produces the output signal, the amplitude can be adjusted to optimise cooling.

[1] To borrow an analogy from Wikipedia: by continuously varying the frequency of our clock, we can continuously adjust it to a second clock by taking the time difference (phase difference) and raising/lowering the frequency of our clock accordingly to minimise this difference. As this is done continuously, both clocks will remain synchronized in time (phase) and therefore also necessarily in frequency [51]

[2] The frequency range for which this is true is termed "Pull in range" and its characterisation is a highly technical topic in engineering, that is far outside the scope of this thesis [52, 53]

For the implementation in the experiment, a Red Pitaya FPGA is used to cool the axial motion, while a Zurich Instruments lock-in amplifier provides the cooling signal for the radial modes.

Digital Bandpass Filter Cooling

For improved cooling performance, especially at lower pressures, a digital bandpass filter implemented on the Red Pitaya FPGA is used. This filter allows for the selective amplification of the feedback signal near the particle's oscillation frequency while suppressing noise at other frequencies. As it is a digital filter, it can also set the phase shift of the feedback signal to make it proportional to the velocity of the particle, similar to the PLL method.
By tuning the bandwidth, it is possible to only take a narrow range of frequencies around the particle's oscillation frequency to feed them back as a cooling signal. This is more robust as the bandpass filter applies to a fixed set centre frequency, and therefore in contrast to the PLL does not need to lock to this signal. Additionally, in contrast to using analogue filters, the digital implementation is experimentally convenient as it allows for tuning the filter parameters while cooling the particle. A gain can also be set accordingly, making it easy to adjust the cooling strength to the particle charge.
As this digital bandpass filter is already running on an FPGA, it can also be easily combined with a PID controller, that sets the output amplitude of the signal. This is useful, as it can be used to adaptively control the gain when the particle motion is excited after a free evolution and decrease it when the particle is cooled down again to the steady cooled state.

3.2. Control equipment

Although the main components of the experimental setup are the trap, electrodes, and detection systems, a variety of control equipment is essential for operating the experiment. This includes voltage sources for applying electric fields, function generators for running the experiment, and data acquisition systems for recording the time traces of the detectors. Each of these components plays a critical role in ensuring precise control over the experimental parameters and accurate data collection, but they are rather standard and not unique to this setup. Nevertheless, we will briefly describe the most important pieces of control equipment used in this experiment, also because we rely on some of their specific properties for the methods developed in this thesis.

DC-Voltage sources

To apply well-defined voltages to the electrodes surrounding the particle, the bias DC voltage comes from Stanford Research Systems DC205 voltage supplies. These provide low-noise, high-precision voltages in range settings of 1 V, 10 V and 100 V. These supplies are chosen for their low noise performance, which is crucial for minimising unwanted

fluctuations in the electric fields that could affect the particle's motion. They also have a built-in voltage ramping feature, which was used as a trigger signal for the voltage scans described in Section 4.3 and Section 4.6, as the falling edge of the busy signal of the supply.
To combine the DC voltages with the high-frequency feedback signals, a custom-built bias tee is used. This allows for the simultaneous application of the DC compensation voltages and the feedback signals. The feedback signals are superimposed onto the DC voltages after they have passed through Mini-Circuits switches, which allow one to turn the feedback on and off quickly with a TTL pulse from the Keysight arbitrary waveform generator.
These voltages are then amplified by Falco amplifiers to reach the required voltage levels for effective compensation and feedback forces in the x and y directions. For the z direction, the voltages from the supply are used directly, as the required voltage levels are lower due to the mentioned force transduction difference in Section 3.1.4.

Function generators

The main function generator used in this experiment is a Keysight arbitrary waveform generator (AWG). This device is used to provide the TTL signals for switching the feedback on and off via the Mini-Circuits switches described above. The model chosen has good enough ($\approx$40 ps) timing precision, which is crucial for synchronising the feedback signals with the AOM switch times. As it also has two channels, it can provide independent control signals for the feedback signals and the laser AOM switching on one device, which saves on issues that could arise when synchronising two different AWGs.

For the experiments with more intricate pulse sequences, an additional Agilent AWG is employed and synced with the Keysight AWG through the busy-out channel. This is only ever used for switching the feedback on and off and therefore has less stringent requirements on the timing precision.

Control Code

The experiment is controlled via a custom Python script with a graphical user interface (GUI) that then interfaces and coordinates with the above devices via USB. To facilitate the measurements in this thesis, the existing GUI controller had to be adapted. Additional features that were added for this thesis include the ability to set voltages after calculating them from the compensation methods described in Section 4.4, safety checks to make sure these amplified voltages are within the permissible range for the amplifiers, a useful interface to also control the AWGs and change their settings via software, the ability to run intricate pulsed sequences for the later experiments, and a more robust data acquisition logic that can detect missed triggers and repeat measurements if necessary. A screenshot of the GUI in its most advanced state is shown in Figure 3.3. The code also received a rewrite to run in a different thread to make sure the GUI does not freeze while running measurements, which is especially important for the later experiments with long

free evolutions, where the user needs to be able to stop the measurement at any time, as the timescales are long enough that the lock or the split levels drift. Additionally, these measurements are rather sensitive, so the ability to stop if things are going on in the shared lab space is very useful to avoid losing the particle and having to start over with pump-down and trapping again. These multiple threads then all print to a common console in the GUI, so the terminal can be minimised and the user can gather all information from the GUI itself.
A measurement then can work in the following way: The settings are configured for all used devices, which are then triggered by the Python script to run the measurement. To ensure concurrency (or rather the correct timing) of the devices, a trigger from one of the interfaced devices is then used to trigger the other devices. For example, for the voltage scans, the Keysight AWG is triggered by the falling edge of the busy signal of the voltage supplies which then triggers the AOM switching and feedback on and off, and triggers the data acquisition on the oscilloscope. This way, all devices are synchronised and the data is correctly recorded for each free evolution of the particle.

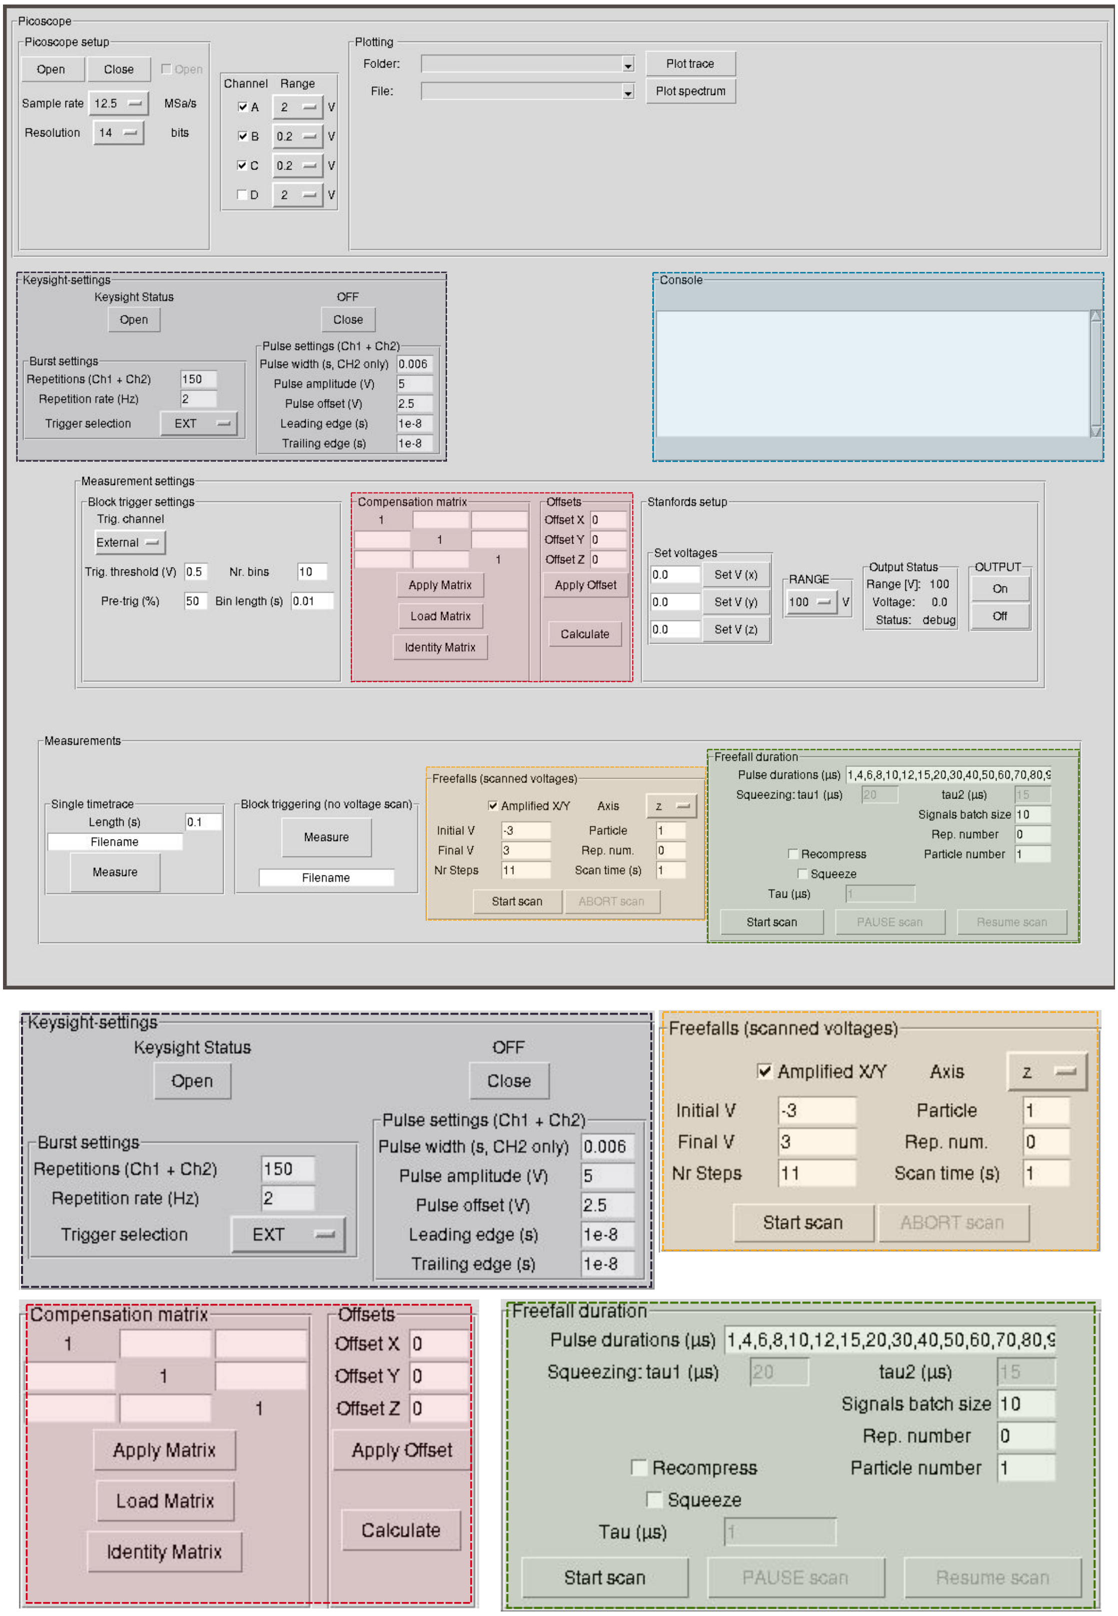

Figure 3.3.: (caption on next page)

Figure 3.3.: **A screenshot of the GUI in its most advanced state.** The coloured boxes indicate the sections that were either added or received a major update for the purposes of this thesis. The black box ("Keysight-settings") is the section for controlling the AWGs, enabling the settings of e.g. repetitions, pulse edges and also pulse width. In the blue box ("Console") is the common console for all threads to print to, which allows the user to gather all information from the GUI itself. The red box ("Compensation matrix and Offsets") is where the compensation matrix can be set, saved and loaded, which can be calculated from the methods described in Section 4.3 and then easily set here. It also includes the ability to set voltage offsets for the compensation, which is useful for the 3D force compensation method described in Section 4.6. The yellow box ("Freefalls (scanned voltages") is where the settings for the voltage scans can be set. This includes the voltage range and steps, as well as the current number of repetitions and the axis that should be scanned. A checkmark indicates whether amplifiers are connected and the voltages are then calculated with the compensation matrix and the amplification factor. The green box ("Freefall duration") is where the duration of the free evolutions can be set, which is used for all measurements where timings are varied. It accepts the list of various timings, pulse duration settings depending on the measurement chosen via checkboxes, and also the number of repetitions of the whole measurement. Additionally, there are a start, pause and resume button.

4. Methods and results

For visualisation, each measurement will have a pulse sequence figure, showing the timing of the different components (trap on/off, cooling on/off, DC bias on/off, Bias, etc). Where relevant, phase space illustrations will be included to visualise the effect of each step on the particle's motional state. All of these phase space figures are illustrative and not to scale. Time is indicated with shading, increasing with the alpha value. All particles here are 156 nm in diameter, SiO_2 nanoparticles, with a mass of approximately 4×10^{-18} kg. The description of experiments is ordered such that hopefully the arising questions are consecutively answered in the following methods, also roughly reflecting the chronological order of experimental development. We will first touch on reheating measurements, which are the most straightforward way to assess the noise in the system and the impact of the DC voltage supplies. It is however a measurement which only captures the energy increase while the particle is trapped, but not the coherence of the state, which is why we will then also perform a recompression measurement to assess the coherence. Having established that the state can remain coherent during free evolutions, we can move on to the compensation methods, which are crucial for enabling long free evolutions in the first place. We will first address compensation voltage scans, which are done to find the optimal compensation voltages, but then show that due to crosstalk between the different directions, this method is not as trivial to extend to 3D as it seems at first glance and a more complex method is needed to find the optimal compensation voltages in all axes. We then introduce a way to measure and correct for this crosstalk. This in turn enables long free evolutions, and by that large state sizes, the main results of this thesis, which also allows us to perform some additional measurements. We will end this chapter with measurements estimating the residual force, and some trial runs to see if squeezing can be achieved with the currently possible free evolution times, which theoretical simulations suggest could be possible in the current setup and represent a promising direction for future work.

Standard methods such as detector calibrations, data analysis methods and some other standard and basic measurements are covered in Chapter A of the appendix Each measurement is split in three parts. First, the method is introduced and the pulse sequence is explained. Then, the measurement is performed, analysed, and finally, the conclusions are drawn. These three parts are highlighted in different colours in the margin to make it easier to follow the flow of each measurement. The colours are consistent across all measurements, so that the introduction, analysis and conclusion of each measurement can be easily identified by their colour for reference.

4.1. Reheating measurements

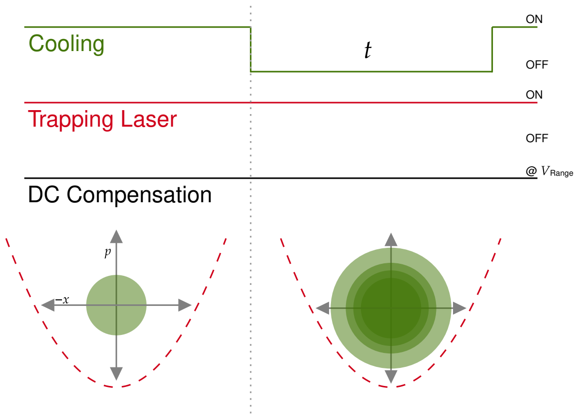

Figure 4.1.: **Pulse sequence for reheating measurements.** For reheating measurements, the particle is first cooled to a low energy state using feedback cooling. Then, the feedback cooling is disabled for a time t of 50 ns, allowing the particle to gain energy due to all of the noise sources that were discussed in Equation (2.27). The measurement runs continuously during the reheating period. After this reheating period, the feedback cooling is re-enabled. The voltage is kept at different levels and range settings, to assess the impact of voltage noise on the reheating rate.

Reheating measurements are particularly important to assess how much noise various components add to the system. In our case, we will particularly focus on the DC voltage supplies and whether there is a significant contribution depending on the specified range of the supply. With this reheating measurement we can measure the combined heating rate including recoil from the trapping laser as outlined in Equation (2.27). This measurement corresponds to a collective measurement of Γ_{bias}, Γ_{recoil} and Γ_{thermal}.

A reheating measurement involves disabling feedback cooling for a time t and measuring the particle's energy increase due to various noise sources. This procedure is sometimes also called a *ring-up* measurement, as the particle's motion timetrace gains amplitude during it and therefore rings-up, compared to a type of measurement where perturbations to a system are measured as they ring back down to the steady cooled state. This is achieved by continuously measuring the variance of the particle's position as a function of time after disabling the cooling, closely following [41].

Reheating measurements at various DC bias levels are shown in Figure 4.2. The measurement consists of disabling the z-feedback for 50 ms and measuring the growth in the variance of the time trace using a moving variance (see Section A.2.1) with a window of 100 datapoints, corresponding to 40 µs due to the sampling rate of 2.5 MHz. Each of these runs is then averaged over the 2250 repetitions, taken in batches of 750 repetitions

each (4500 repetitions for the reference measurement with DC bias disconnected, one before and one after varying the DC-Voltages) at pressures below 6×10^{-8} mbar and then used to compare the different bias settings. To these averaged variance increases we also fit a linear function to extract the slope and then plot these in Figure 4.3.
We can see that voltage noise only becomes significant (compared to the batch with the DC supplies disconnected) when the axial voltages are held above 50 V, over an order of magnitude beyond typical compensation voltages. Furthermore, the measurements with 10 V axial bias were performed using the 100 V output range of the DC supply, for which the RMS noise is significantly higher than in the 10 V or 1 V range. This implies that the increased heating rate at 75 V and 100 V is due to the relative stability of the supply (specified as 1 ppm for 24 h) rather than an increase in RMS noise due to a larger voltage range. From this we note that this is clear evidence that while physically, any environment charge state could be compensated for, the practical limitations of the voltage supplies mean that there is a limit to how much compensation can be achieved before voltage noise becomes a significant contributor to the reheating rate. As it is a reheating *rate*, this becomes more significant for longer total measurement times without cooling.

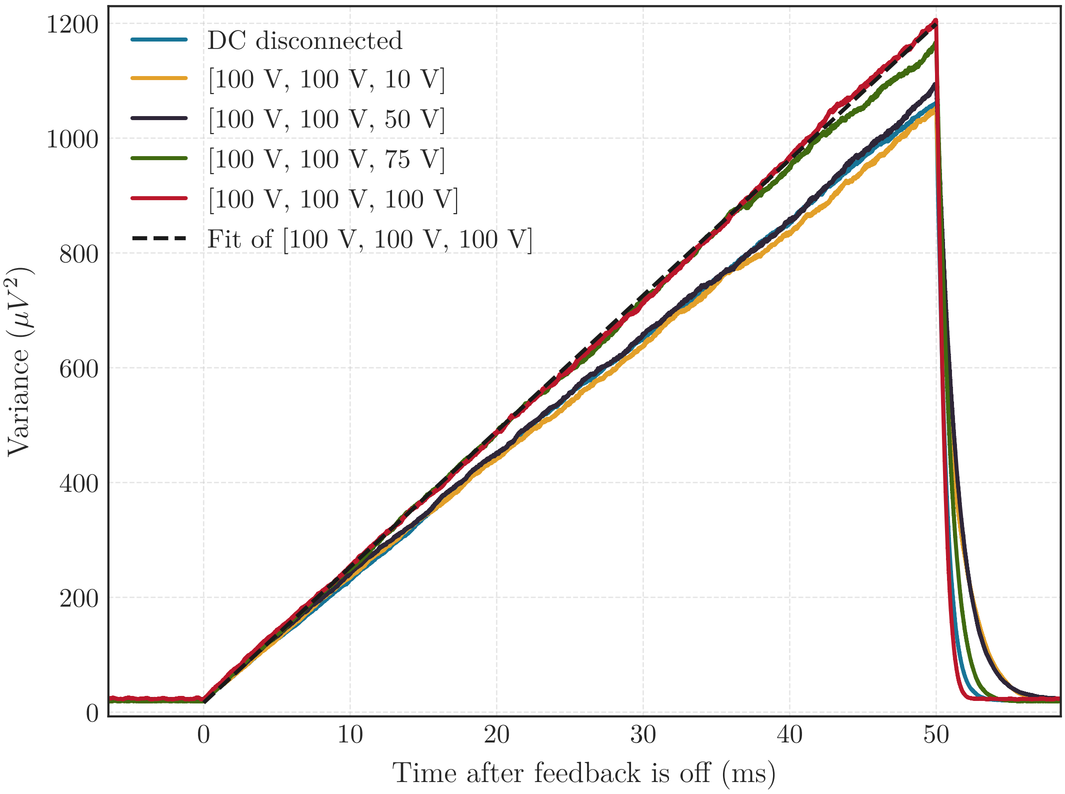

Figure 4.2.: **Reheating measurements at various DC bias levels.** Each curve shows the average increase in variance of the position detection over multiple runs at different DC voltage levels. The voltage range on the supply was set to 100 V for all axes and measurements, the actual voltage set to be output at the corresponding label $[V_x, V_y, V_z]$. Each curve is the average of 2250 repetitions (4500 repetitions for the reference measurement with DC bias disconnected, taken before and after varying the DC-Voltages, to be able to distinguish from drifts.)

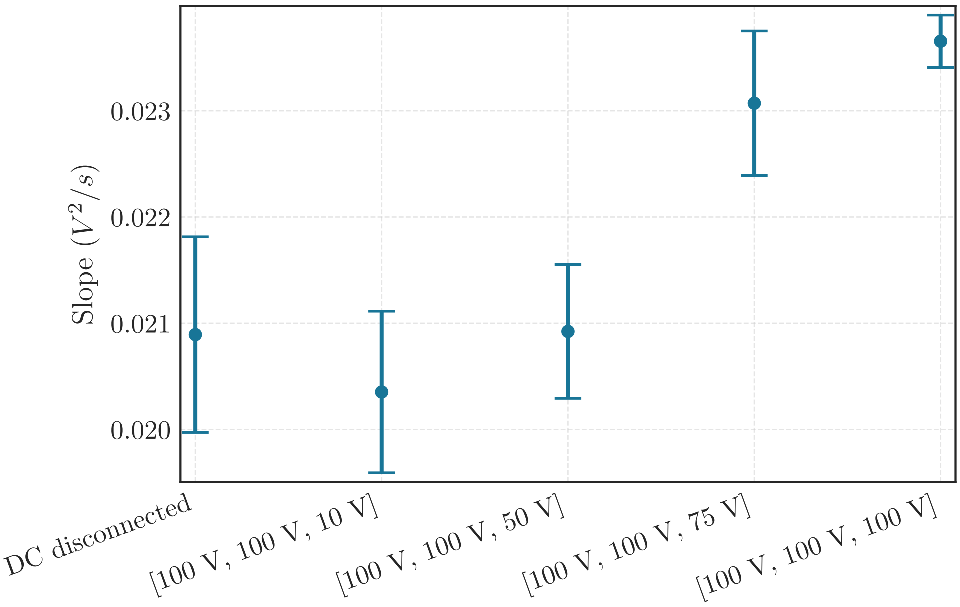

Figure 4.3.: **Uncalibrated reheating rates at various DC bias levels.** The slopes of the reheating curves in Figure 4.2 are extracted by fitting a line to the linear part of the curve, which is between 0 ms and 50 ms. The resulting slopes are then plotted as a function of the DC bias level, again as in Figure 4.2 with the labels indicating the set bias in the three directions as $[V_x, V_y, V_z]$. The plot shows that the reheating rate remains close to the rate at disconnected DC supplies for bias levels that are typical for measurements (z-bias $\leq 10\,\mathrm{V}$), but increases significantly at higher bias levels, particularly above 50 V. This indicates that voltage noise becomes only a dominant contributor to the reheating rate at these high bias levels. The error bars show the standard deviation of the slope fit parameter.

4.2. Coherence Measurements

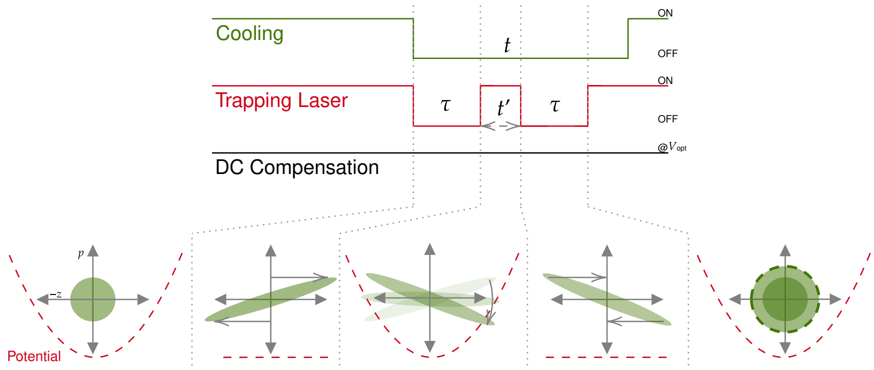

Figure 4.4.: **Pulse sequence for recompression measurements.** The particle is released from a cooled state for a time τ, during which it shears as it evolves freely. A short pulse for time t' of the harmonic trapping potential rotates the state in phase space, such that a second free evolution for the same time τ recompresses the state to its initial extent if coherent. The particle's motion is recorded after the second free evolution to assess the degree of recompression, which is indicative of the coherence of the state during the free evolutions, a ratio of 1 indicating perfect recompression and therefore coherence. The cooling is off in all axes for 6 ms, the DC-bias is on at the optimal voltages during the whole sequence.

State expansions for nanoparticle interferometry protocols need to be coherent to be useful for state preparation of large delocalised states, as a reduction in purity would wash out the very same features we try to produce.
Having looked at the reheating measurements in Section 4.1, which capture the reheating rate consisting of Γ_{recoil} (photons) and Γ_{thermal} (gas), we notice this does not capture the heating rates from noise in the feedback loop (a part of Γ_{feedback}). Also not included is any contribution due to shot-to-shot dephasing, as is for example seen in Larmor spins [54] which arises from slight mismatches in frequencies from shot-to-shot.
In this measurement, decoherence manifests in a larger state after the protocol run. The free evolution itself also changes the state, so we would prefer a protocol that ends in a symmetric, circular state that is easily comparable with the initial state. If they are the same, this implies no significant loss of purity during short (5 µs) free evolutions with 3D compensation. To assess this, we perform a recompression protocol.
The recompression works as follows: We start with a cooled particle, cooling switched off in all axes for 6 ms with the first laser pulse. This is where we switch off the trap and let the particle evolve freely for a time $\tau = 5\,\mu\text{s}$. After this free evolution, we the trap is switched on for a time t', followed by another free evolution for time τ.
Physically, this can be understood as first squeezing the state during the first free evolution, which then after τ is squeezed along an angle θ in phase space (see Section 2.2

and Equation (2.17)). The laser turned back on for time t' then applies a rotation in phase space (it is a harmonic potential), which can be chosen to be 2θ such that after the second free evolution for time τ, the state is recompressed to its initial extent if no entropy is added to the system. To see this behaviour, t' is scanned around the optimal value. This pulse sequence and phase space is illustrated in Figure 4.4, the measurement result is shown in Figure 4.5. A return to the initial variance indicates no increase in entropy.

An important note on Figure 4.4 is that any integer 2π rotation will recompress but add noise from photon recoil due to the longer pulse duration in which the light is on. While not ideal, also the first minimum has its difficulties as it requires a very short pulse duration, recompressing for pulse times of 1 µs. As a compromise, we focus on the second minimum.

Each timetrace is sectioned in pre- and post-free-evolution parts. We then do multiple fitting passes (see Section A.2.2) to extract the amplitudes, frequencies and phases of the three axes of the oscillation in both sections of the time trace. From the fitted parameters we can reconstruct the sinusoidal trajectories, discarding the slope and offset parameters as these are not relevant for the variance of the state. We calculate the variance as a function of time and then find the maximum and minimum variance during the post free evolution part of the timetrace. As we scan t', there is not a single post-evolution variance, since for any non-ideal value of t', the state will not be perfectly recompressed and still asymmetric. Using the maximum variance, we can make guarantee we are at worst under-estimating our recompression (see Section A.2.3).

We also use a basic Lyapunov differential equation (as introduced in Section 2.2) approach to model the expected variance evolution as a function of the pulse duration t', to get some intuition about the expected behaviour. It is important to note that this is 1) a very basic model with very rudimentary checks on the assumption and 2) the main data of note is just the point of maximum recompression which comes directly from the recorded data.

In Figure 4.5 we can see that if we start with too little rotation time t' , the post-evolution position variance is larger than the initial extent, reaching almost 1 within the standard deviation for the optimal t' and increasing for rotations further than the ideal 2θ. The best recompression we reach is (1.12 ± 0.07), but this is not for the first possible minimum and also not for the best timing of t'. We also see that it agrees well with our basic theory model. There is a subtlety in this measurement and the interpretation thereof: recompression without state tomography is technically a necessary but not sufficient condition for coherence [1]. Nevertheless we would see added force noise (a change in the

[1] The argument being, that purity is given as $(\det \boldsymbol{\sigma})^{-\frac{1}{2}}$, with $\boldsymbol{\sigma}$ here standing for the covariance matrix of the Gaussian state [44]. Without state tomography, we cannot obtain σ and therefore cannot make a definite statement on the overall loss of purity during the protocol. The protocol is sensitive to any decoherence that inflates the covariance matrix, but is blind to decoherence that preserves the marginal position and momentum variances. Intuitively, this can be explained as the possibility of producing, through *some* unaccounted for effect (nonlinear dynamics, etc.), a state which is not the initial state but has the same variance marginal in position. As we do not see the full state, we could miss this in our analysis.

momentum quadrature during the protocol) as well as position noise (a change in the position quadrature), which we measure to not be significantly present. Some other work has come around this issue by doing retrodiction, which presents a solution by–simply speaking– calculating backwards from the recapture timetrace a full state at the end of the second free evolution, giving error bars on all state parameters, including deviations from gaussianity [8].

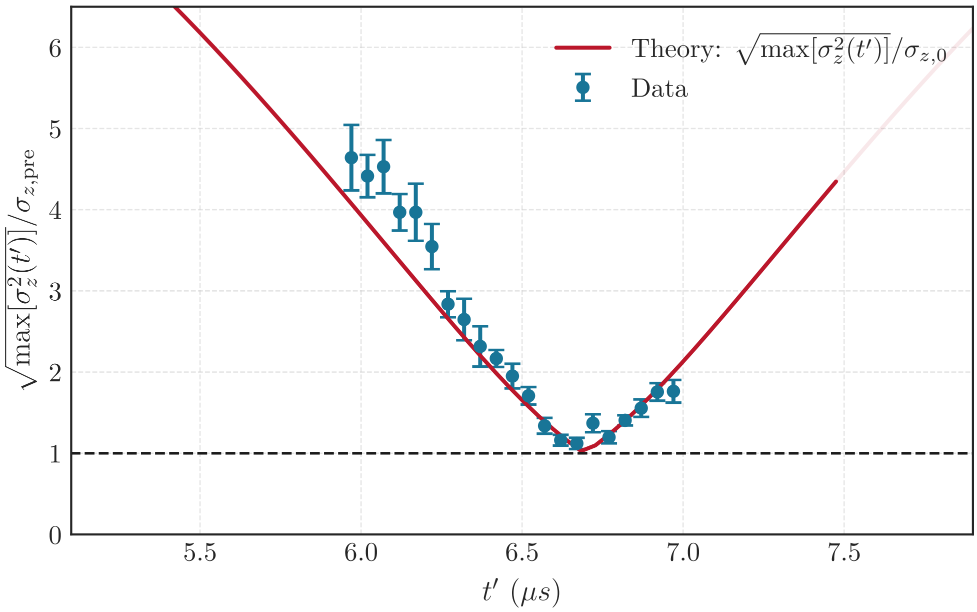

Figure 4.5.: **Recompression protocol, maximum variance as a function of t'.** State trajectories are reconstructed by fitting pre- and post-free-evolution time traces to extract oscillation parameters in all three axes. To conservatively estimate recompression across scanned pulse durations t', the maximum post-evolution variance of the reconstructed oscillations is used. At the optimal t', the variance ratio recovers nearly to unity (1.12 ± 0.07), demonstrating recompression close to the initial spatial extent (taking into account photon recoil, non-optimal timing choice, see main text). This recovery agrees with a basic Lyapunov variance model and establishes an upper bound on decoherence, fulfilling a necessary condition for state coherence. The best fit with the model is achieved for a small offset of $-0.03\,\mu s$ which is well explained by a small difference in the AOM rise and fall times.

4.3. Free evolution scans for compensation voltages in one direction

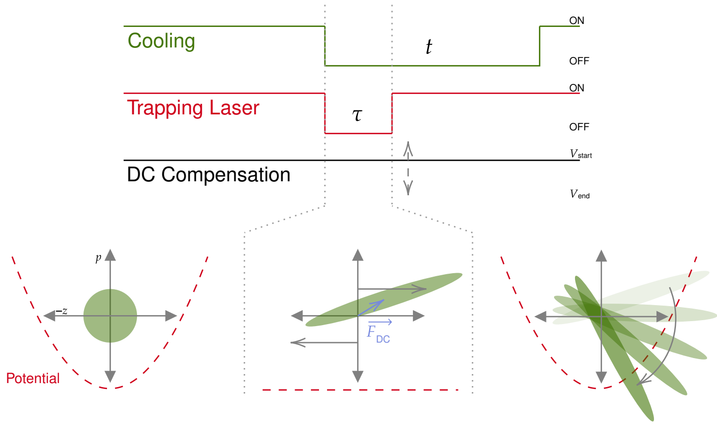

Figure 4.6.: **Pulse sequence for compensation voltage scans.** First, the particle is cooled, then released for a time τ. This is done `nBin` times at a certain compensation voltage to get statistics. Then the compensation voltage is changed to the next value, scanning a range of voltages. The particle's motion is recorded before releasing it and after recapturing, allowing to extract the voltage for which the energy increases the least after free evolution, which is then the optimal compensation voltage in that direction.

To compensate for stray fields, we do not need to know the absolute values of the charge of the particle or the environment. We can use the fact that a static force leads to displacement and velocity increase in the trap and therefore to an increase in the particle's mean energy after a free evolution (see Equation (2.23)). By scanning the compensation voltage in one direction while performing free evolution measurements at each voltage, we can find the optimal compensation voltage that minimises the energy increase after the free evolution. This increase is not necessarily zero due to other noise sources, but the minimum of the curve provides the best compensation voltage in that direction, and thus the best achievable compensation. Noteworthy, although straightforward, is the fact that this environmental charge compensation is independent of the particle's charge[2]

The procedure is as follows: First, a compensation voltage is set in one direction (e.g., z), and a series of free evolution measurements are performed (to gather statistics on each voltage and avoid being affected by variations of each single run) with a fixed free evolution time τ (e.g. 10 µs). This free evolution time is chosen to be long enough to clearly observe the effect of the quadratic force scaling mentioned in the theory section

[2]Since the compensation force $F_{\text{comp}} \propto q_{\text{part}}$ and the force from the environment $F_{\text{env} \propto q_{\text{env}} q_{\text{part}}}$, as we reach the optimal compensation voltage when $F_{\text{comp}} = F_{\text{env}}$, the particle charge q_{part} cancels out, only leaving the environment charge q_{env}.

(Equation (2.23)), but short enough not to make the particle's motion dangerously excited even if the environmental charge is high. If there is no clear parabola visible, the steepness scales with $\tau^2 + \tau^4$ (see Equation (2.23) looking at the constant F_z), so increasing τ will make the parabola steeper and easier to identify.
For each free evolution measurement, the particle's motion is recorded before and after the free evolution to determine the mean energy increase, as illustrated in Figure 4.7. This is repeated for a range of compensation voltages around an initial guess. To get an estimate of the error, this is done five times for the same voltage. Each timetrace is split up into pre and post free evolution time segments to remove the free evolution part including a buffer time of 3 µs after recapture for detector transients as there is no detection and light available during the free evolution time. As the cooling is only reenabled after data acquisition, any heating during this buffer does not contribute to the resulting V_{opt} as it is a constant offset to all points. Afterwards, both the pre and post free evolution time traces are turned into power spectral densities (PSDs) which are then integrated around the respective frequencies to extract the energy in each mode before and after the free evolution, to get the increase in the mode of interest. The exact method is described in Section A.2.6. The mean energy increase is then plotted as a function of the compensation voltage, and a quadratic fit is used to interpolate the minimum increase and therefore the best compensation voltage. The fitting procedure is further described in Section A.2.4 of the appendix.

In Figure 4.7 we can see that the measured increase in mean energy does indeed show parabolic behaviour as a function of the compensation voltage, with a clear minimum at (0.33 ± 0.02) V The insets show the time traces analysed for the variance extraction, showing a lower amplitude after a voltage closer to the fitted optimal voltage, as well as an expected phase shift of the ensemble of time traces when going through the optimal point, indicating a shift of the particle through the potential minimum.

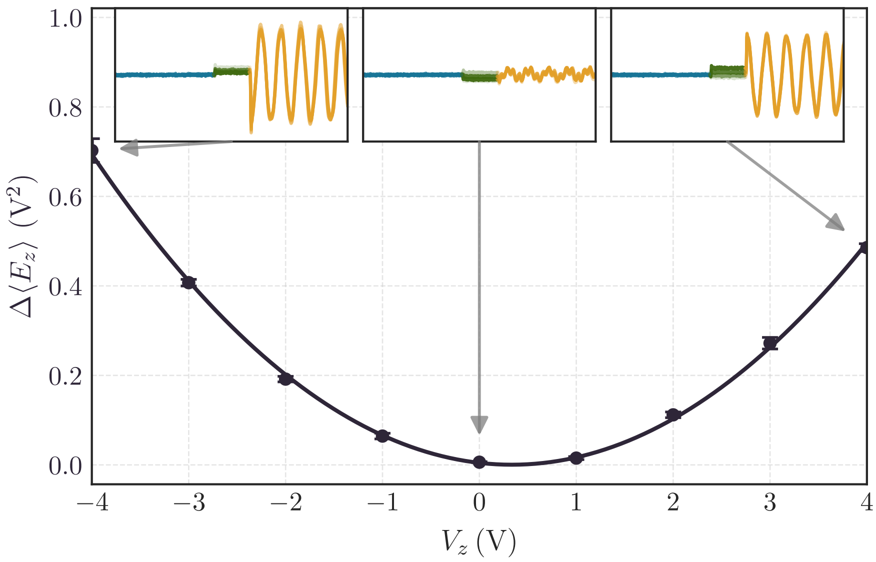

Figure 4.7.: **Example data of a compensation voltage scan in one direction (here z).** Each point represents the mean energy increase after free evolution for a given compensation voltage, with error bars showing the standard error of the mean of five free evolutions at that voltage. The solid line is a quadratic fit to the data, with the minimum indicating the optimal compensation voltage in this direction. **Insets:** Example time traces of the five free evolutions overlaid on top of each other at the corresponding voltage. It can be seen that the oscillation amplitude after free evolution is minimised near the optimal compensation voltage, as well as the expected phase shift of the ensemble of time traces when going through the optimal point, indicating a shift of the particle through the potential minimum. The fitting implementation is included in Section A.2.4, so is the PSD integration method for the energy extraction in Section A.2.6.

4.3.1. Large-range z scan and nonlinearities

Extending the compensation voltage scan to its maximum, assuming robust cooling upon recapture to manage even a particle with relatively high motional energy that no longer solely explores the harmonic part of the potential, this method can also be used to map out the trap's nonlinearities. This knowledge can then provide estimates on the required compensation accuracy for a given free evolution duration τ. This is illustrated in Figure 4.8, where a large-range scan along z is shown, with the particle exploring the non-harmonic region. In this region, the harmonic approximation from Equation (2.10) breaks down, also meaning that the directions of motion can no longer be treated independently, and the particle's motion becomes more complex. With knowledge of the trap waists in each direction, this can also be converted to a required accuracy in the radial direction.
To find the required accuracy, these one directional scans in z with a detector calibration from Section A.1 in conjunction with force calibration can be used to convert the voltage signal to a position, showing that nonlinearities start at $\approx 170\,\mathrm{nm}$ displacement along z. This stems from the fact that the particle is released for a time τ with a calibration factor α, and the classical displacement due to a static force F which can be reformulated as:

$$z = C_{\mathrm{NV}}(V - V_{\mathrm{opt}})\frac{\tau^2}{2m}. \tag{4.1}$$

The details of the force calibration are briefly described in Section A.5 of the appendix.
Since the trap waists are different in the axial and radial directions, we can estimate the required accuracy in the radial direction by using the ratio of the waists, which is approximately 1:3, meaning that the required voltage accuracy to compensate for a 100 µs free evolution in the radial direction is approximately 20 V while in the axial direction it is approximately 1 V. Going towards the 1 ms regime, the required accuracy in the radial and axial direction becomes even more stringent, on the order of 0.2 V and 0.01 V respectively.

4.4. Cross Talk between directions

As discussed in the previous section, scanning across a voltage is an excellent way to find the optimal voltage in one direction. It would then lend itself to simply extend this and to do this iteratively in all three directions until three voltages were found with enough confidence and accuracy to cancel the forces in 3D and enable long free evolutions. However, a mode analysis, integrating the PSD around the respective frequencies, and looking at two modes while performing a z-scan shows this is not the case. In Figure 4.9, the dashed red line shows the behaviour of the x mode. If it were independent of the z-compensation voltage, it would be flat. However, it shows a clear dependence on the z-compensation voltage. The same is true for the y mode (not shown in the figure). The analysis was done the same way as in Section 4.3.

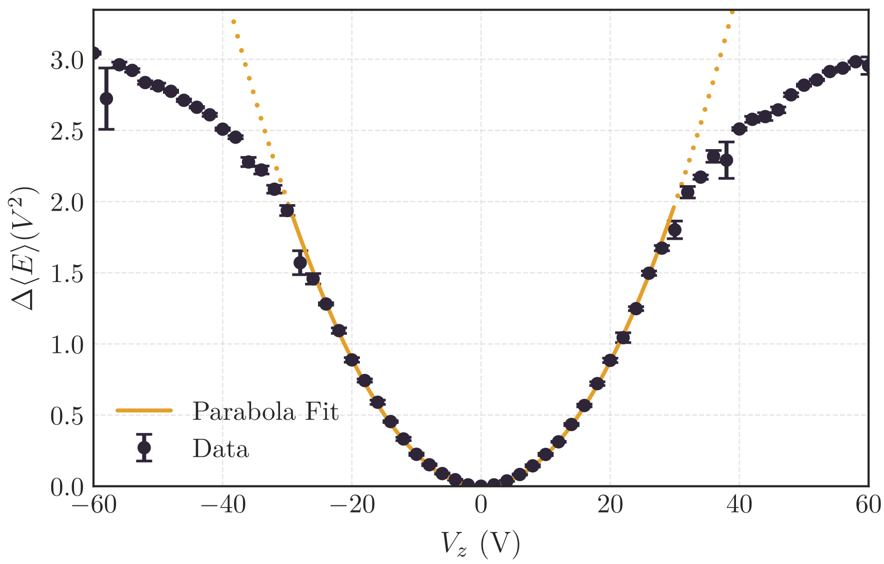

Figure 4.8.: **Large-range z-compensation voltage scan showing the deviation from the harmonic potential at large voltages.** The points show the increase of the variance after a free evolution of $\tau = 15\,\mu\text{s}$ for different z-compensation voltages, with error bars showing the standard error of the mean of five free evolutions at each voltage. The yellow solid and dotted line is a fit to the data using the function $a(V - V_{\text{opt}})^2 + b$, where a, V_{opt} and b are fitting parameters. The solid part of the fitted line indicates which data points are used for the fit. The fitting implementation is included in Section A.2.4. The deviation from the quadratic behaviour at large voltages shows the onset of nonlinearities in the trap potential and the detection which become significant when the particle explores regions far from the trap centre.

This crosstalk stems from the fact that we apply the voltage to the electrodes but observe the motion in the trap frame. These two frames are not aligned, and therefore a method to measure the transformation matrix between them is needed for proper compensation.

In general we can separate such a transformation into a scale matrix S and a crosstalk matrix $\tilde{C}^{-1}$, where the scale matrix contains the three scale factors that are needed to convert the voltage applied to the electrodes to a force in the trap frame, and the crosstalk matrix contains the relative contributions of each electrode to each trap axis. One matrix scales our transformation to the correct magnitude, while the other one ensures that the force is applied along the correct direction.

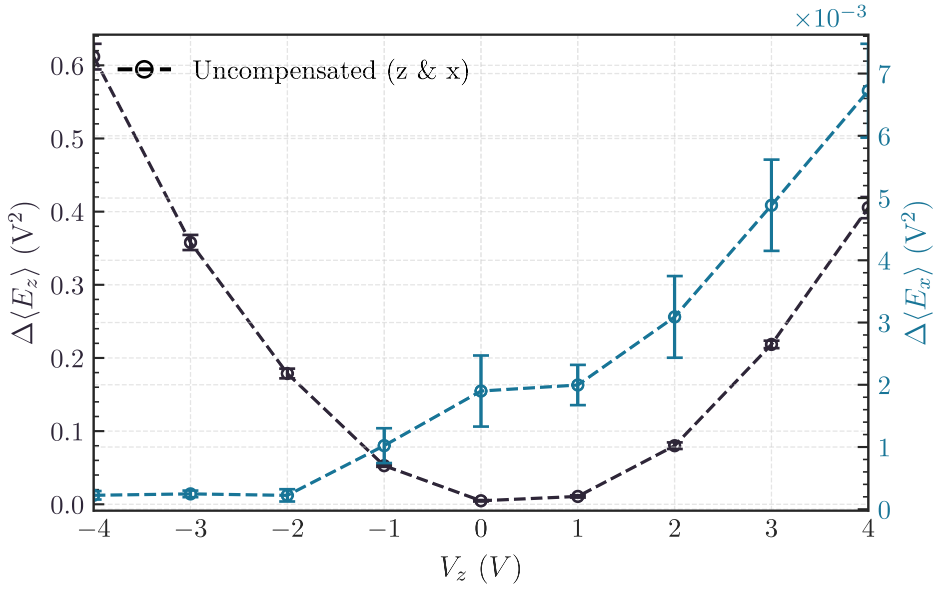

Figure 4.9.: **Force crosstalk between different directions during compensation voltage scans.** We compute power spectral densities using 5 ms of data before and after the 10 µs free evolutions to isolate the response of the individual motional modes as done in Section 4.3. The z mode shows the expected parabolic behaviour (black), while the x mode shows a clear dependence on V_z due to cross-talk (blue dashed line).

These scale factors are not necessarily needed (although they could be obtained) as we are not interested in the correct magnitude of the applied forces. For the purpose here, it is already enough that the force applied is well defined within the trap axes; a slight shift in magnitude is acceptable, as all values are found empirically by scanning and applying voltages. We measure the least excitation in the modes after applying a certain voltage, then apply that same voltage to compensate the field. What precisely of this voltage reaches the particle, is insubstantial to the compensation. We can therefore rewrite the matrix as:

$$C^{-1} = \begin{pmatrix} 1 & \frac{V^x_{x'}}{V^x_{y'}} & \frac{V^x_{x'}}{V^x_{z'}} \\ \frac{V^y_{y'}}{V^y_{x'}} & 1 & \frac{V^y_{y'}}{V^y_{z'}} \\ \frac{V^z_{z'}}{V^z_{x'}} & \frac{V^z_{z'}}{V^z_{y'}} & 1 \end{pmatrix} \begin{pmatrix} \frac{1}{C_{x'x}} & 0 & 0 \\ 0 & \frac{1}{C_{y'y}} & 0 \\ 0 & 0 & \frac{1}{C_{z'z}} \end{pmatrix} := \tilde{C}^{-1} S. \tag{4.2}$$

4.5. Measuring the Cross-talk Correction Coefficients

To address this experimentally, we use a method we call *cross-cooling*. We will first examine the drawbacks of the simpler method: driving the particle with a sinusoidal signal and observing the response in the different modes. The problem with this method is that the particle's response is determined not only by the driving force but also by its susceptibility and thus the detuning to the particle's resonance frequency at that moment in time. The resonance frequency is not perfectly stable but fluctuates in the 1 kHz range, which makes the drive response amplitude also dependent on these fluctuations and therefore difficult to measure accurately. It is especially critical that one finds a detuning small enough to be picked up in all modes (emanating from all electrodes), but also large enough to first disentangle it from the particle's resonance and second, not to excessively heat/drive the particle.

The solution to this problem is to use feedback cooling to measure the crosstalk. The idea is to cool one mode (e.g., x as in Figure 4.10) with a certain amplitude on the corresponding electrode (x), then observe the demodulated signal amplitude of the mode and record it. Afterwards, the same mode (z) is cooled, but now with a different amplitude (and potentially a 180° phase-shift) on another electrode (e.g., y). The demodulated signal amplitude of the z mode is then used to vary the amplitude of the feedback signal until the same demodulated amplitude as before is reached. The ratio of the two amplitudes then provides the relative crosstalk from the y-electrode to the z mode. Repeating this for all combinations of modes and electrodes then yields the full transformation matrix C^{-1} from the electrode frame to the trap frame, as each iteration yields an equation of the form $C_{x'i}V_{ix'} = C_{y'i}V_{iy} = C_{z'i}V_{iz'}$ for each trap axis $i \in \{x, y, z\}$. These equations are then used to determine the entries of C^{-1} up to the three scale factors given by $C_{i'i}$. To measure the required voltage ratios, we first need to ensure the detections do not have optical crosstalk by manipulating the D-shaped mirrors in the detection paths. This results in a suppression of the other modes in each detection according to Table 4.1. The resulting columns of $\tilde{C}^{-1}$ then define the proportion of voltages needed to apply a force along a given trap axis, thus allowing for iterative compensation scans for each axis. The scale matrix S contains the mentioned scale factors that are not needed for these compensation procedures.

With these optical crosstalk free detections, we can then perform the cross-cooling method to find the required voltage ratios. The measured signal strengths for this procedure are summarised in Table 4.2
The values found for the cross talk matrix $\tilde{C}^{-1}$ in this setup with linear polarisation are:

$$\tilde{C}^{-1} = \begin{pmatrix} 1 & 0.32 & -37 \\ 0.36 & 1 & 4.4 \\ 0.0011 & -0.0012 & 1 \end{pmatrix}. \tag{4.3}$$

It is worth mentioning here, that we found these values to be geometric in origin and therefore constant over long periods (months), but polarisation-dependent, as the trap axis differs for, say, a more circular polarisation.
The precision with which this method removes crosstalk can qualitatively be seen in the mode analysis in Figure 4.11, which now shows the blue solid line of the x mode during a z-compensation scan being flat within the error bars.
With this matrix determined, compensation scans can now be performed to iteratively find the optimal compensation voltages in all three directions.

Table 4.1.: **Measured signal strengths of optical cross talk for determining the crosstalk matrix,** showing the suppression of the other modes in each detection. The values are given in decibels, with the mode being detected in each row indicated in the first column and the mode being observed in each column indicated in the header.

	x mode (dB)	y mode (dB)	z mode (dB)
x detection	21	< 3	3
y detection	< 3	21	3
z detection	30	38	69

4.6. Free evolution scans for Compensation Voltages in three directions

Now with the crosstalk matrix determined, we can perform compensation scans in all three directions iteratively to find the optimal compensation voltages in all three directions.
We use the same pulse sequence as Figure 4.6 but now V_x, V_y, and V_z are applied according to the crosstalk matrix. In the analysis, we use again the mode analysis from Section 4.3.
An example of such 3D scans is shown in Table 4.3 also showing a representative set of the parameters used for each scan and plots for a representative scan.
The general procedure is as follows: First, a longer-range scan in z with a conservative free evolution time is performed to find the optimal z-voltage. Then, using this initial

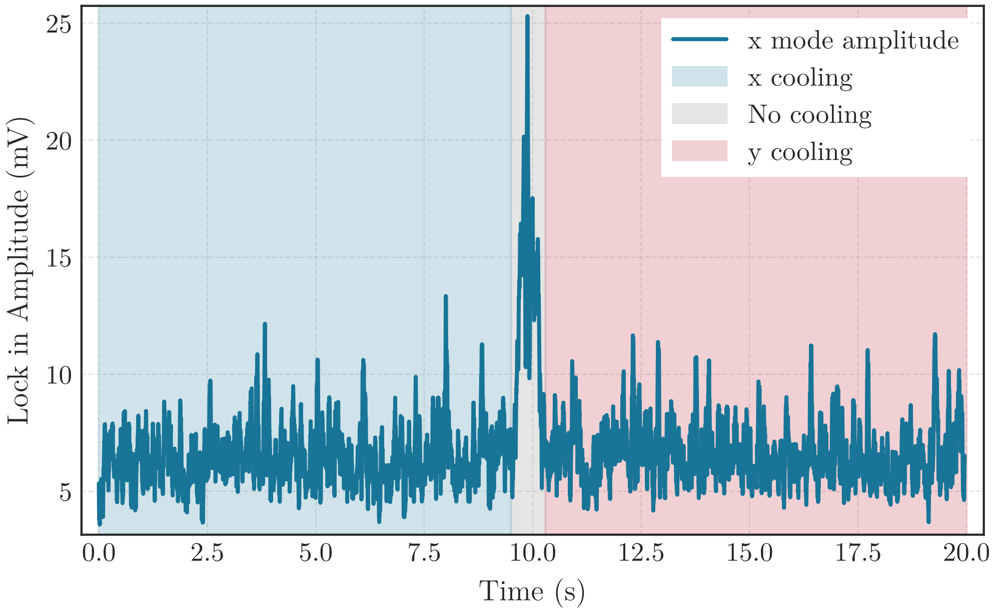

Figure 4.10.: **Demodulation signal of the x detection during one cross cooling measurement.** We see a lock in trace of the x mode at 7×10^{-2} mbar. In the left shaded region, the x mode is cooled with a certain amplitude on the x electrode, and in the right shaded region, the x mode is cooled with a different amplitude on the y electrode, but to the same lock-in amplitude. In this case, a roughly factor 4 increase. The ratio of the two amplitudes then gives the crosstalk from the y electrode to the x mode. The same procedure is repeated for all combinations of modes and electrodes to find the full crosstalk matrix and results in the measured voltage ratios in Table 4.2.

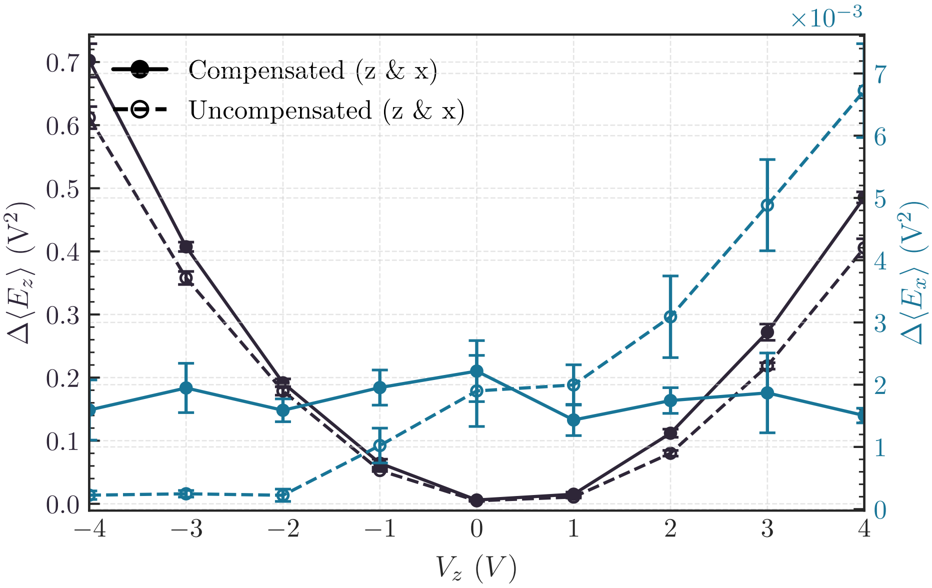

Figure 4.11.: **Force crosstalk (removed) between different directions during compensation voltage scans.** We compute power spectral densities using 5 ms of data before and after the 10 µs free evolutions to isolate the response of the individual motional modes. The z mode shows the expected parabolic behaviour (black), while the x mode shows a clear dependence on V_z due to cross-talk (blue dashed line). When the voltages are instead scanned in a proportion given by the third column of C^{-1} (solid lines), the cross-talk into x is significantly reduced. Each point contains five repetitions and error bars are standard errors of the mean.

Table 4.2.: **Measured voltages ratios for determining the crosstalk matrix $\tilde{C}^{-1}$**, where the first index indicates the mode being cooled and the second index indicates the electrode on which the cooling signal is applied. The values are given in volts and the phase of the feedback signal is also indicated, as it is relevant for calculating the crosstalk matrix.

	x-x (V // phase)	y-x (V // phase)	z-x (V // phase)
x detection	1.5 // 0°	5.53 // 180°	0.039 // 0°
	x-y (V // phase)	y-y (V // phase)	z-y (V // phase)
y detection	4.35 // 160°	1.5 // -20°	0.087 // 160°
	x-z (V // phase)	y-z (V // phase)	z-z (V // phase)
z detection 1	5.9 // 125°	5.9 // -55°	0.01 // -55°

guess for the z-voltage as a bias, a scan in x is performed to find the optimal compensation voltage in x. As the mode is less sensitive and an offset given by an unbiased z-voltage is already taken into account, the free-evolution duration τ is also increased to 20 µs for this and the next scan. For the following y-scan, it is not necessary to bias in x, resulting in more or less the same parameters as for the x-scan, except the range that differs due to the crosstalk matrix. Finally, a last detailed scan in z is performed to refine the z-compensation voltage found in the first scan, now biasing in both x and y with the voltages found in the two respective scans. This last scan has a longer free evolution time to increase sensitivity; a value of around 40 µs often suffices, as the other two directions are already compensated, leading to a more stable measurement.

If any of the scans do not show a clear minimum, the procedure can be repeated with longer free evolution times to increase sensitivity until a clear minimum or parabola is found. For the best possible compensation voltages, also statistics are important and can be gained by performing the same scan multiple times. The resulting standard error and the mean of the optimal voltages are then used to determine the final compensation voltages, in accordance with the needed precision we found in Section 4.3.1. The voltage triplet found in this way for the measurements in Table 4.3 is $[(-9.2 \pm 0.2)\,\mathrm{V}, (6.5 \pm 0.3)\,\mathrm{V}, (1.862 \pm 0.001)\,\mathrm{V}]$ for $[V_x, V_y, V_z]$ respectively. This compensation voltage triplet was then used for the free evolution measurements in Section 4.11.

It is also possible that only the flanks of the parabola are observed, in which case extrapolation is unreliable. However, we explore environment charge manipulation methods in Section 4.7.3 to ideally shift the minimum into the scan range.

Run	Free evolution (μs)	Scan	Start (V)	End (V)	Pressure (mbar)	x bias (V)	y bias (V)	z bias (V)	Notes	Plot Ref.
0	10	z	-3	3	6×10^{-8}	0	0	0		Fig. 4.12a
1	10	z	0	4	6×10^{-8}	0	0	0		Fig. 4.12b
2	20	x	-79	220	6×10^{-8}	0	0	1.88		Fig. 4.12c
3	30	x	-79	50	6×10^{-8}	0	0	1.88		Fig. 4.12d
4	40	x	-75	55	6×10^{-8}	0	0	1.88		Fig. 4.12e
5	20	y	-150	141	6×10^{-8}	-10.66	0	1.88		Fig. 4.12f
6	40	y	-60	60	6×10^{-8}	-10.66	0	1.88		Fig. 4.12g
7	40	z	1.6	2.2	6×10^{-8}	-10.66	8.14	0		Fig. 4.12h
8 - 14	40	x	-75	55	6×10^{-8}	0	8.14	1.8621	Mean: -9.1929 V St. dev.: 0.239 72 V	Fig. 4.12i
15 - 21	40	y	-55	65	6×10^{-8}	-9.1929	0	1.8621	Mean: 6.5243 V St. dev.: 0.306 83 V	Fig. 4.12j
22	40	z	1.4	2.2	6×10^{-8}	-9.1929	6.5243	0		Fig. 4.12k
23-27	40	z	1.35	2.25	6×10^{-8}	-9.1929	6.5243	0	Mean: 1.8582 V St. dev.: 0.002 289 6 V	Fig. 4.12l

Table 4.3.: **3D compensation scans used to locate optimal compensation voltages in all axes.** The scans show the predicted scaling with F^2 in all axes, with a minimum visible in each. The solid line is a fit through the increase in mean energy obtained via mode analysis. The compensation matrix is applied, meaning that a 0 V bias does not mean a zero voltage output on the corresponding power supply. All measurements span 11 different voltages, releasing and recapturing the particle 5 times for each. **Run 0-1** Initial scans in z, long range to find the minimum. Laser is shut off for 10 µs. **Runs 2-4** Scan in x, biased in z with the voltage found in the initial z scans. Laser is shut off for 20 µs,30 µs,40 µs to increase sensitivity. **Runs 5-6** Scan in y, biased in z and x. Laser is shut off for 20 µs and 40 µs. Then a second, zoomed in iteration of all modes starts with 40 µs, where the scans are repeated multiple times to gain statistics on the optimal voltages. The resulting mean and standard deviation of the optimal voltages are then used to determine the final compensation voltages, in accordance with the needed precision we found in Section 4.3.1.

4.7. Charge calibration

There are multiple charges to be considered in the experiment. One of them is the (relative) environmental charge, which we measure all the time by looking at the compensation voltages. More complex is the measurement of the *particle* charge, and this requires a dedicated measurement and calibration procedure, which we will describe in the following.

4.7.1. Particle Charge

For trapped nanoparticles, there are a number of existing techniques for particle charge control and characterisation, including high-voltage discharge [55] the approach used in this thesis. Charge control would be a great tool to have for long evolutions, due to the fact that by having only a lightly charged particle, it would be less susceptible to environmental charges, and all signals applied for compensation could in theory just be scaled in voltage.

Force sensing reveals that plasma discharge significantly charges the environment, showing as an order-of-magnitude increase in the required axial compensation voltage. Consequently, we avoid discharges and utilise natively charged particles (tens of elementary charges) for the measurements. While focused UV illumination could offer a more targeted alternative for charge modification,[56], we chose to restrict the high-voltage discharge method strictly to the initial characterisation of the detector signal-to-charge ratio.
To perform the calibration, a phase-locked loop (PLL) is locked to the particle's axial motion signal at a chosen frequency of 150 kHz, and a resonant oscillating drive voltage of 10 Vpk is applied to the axial electrodes to excite a steady-state oscillation amplitude at this frequency. A high-voltage power supply is then briefly turned on at 1500 V and a current limit of 10.5 mA prevents excessive current flow to the electrode in the vicinity of the particle, producing a plasma discharge that stochastically adds or removes elementary charges from the particle. The signal level of the demodulated lock-in amplitude changes in discrete steps of one or multiple elementary charges each time the particle's charge state shifts. The resulting time trace is shown in Figure 4.13.
First, the raw signal is low-pass filtered by convolution with a normalised Gaussian window of 8000 samples, giving a clean signal to analyse further. Charge-jump transitions are then found by thresholding the absolute value of the discrete first difference of the filtered signal at 4×10^{-8} V. All samples at or above this threshold are masked (set to `NaN`) so that only the flat charge-level plateaus contribute to the subsequent analysis. A 400-bin histogram of the remaining plateau samples is computed, and its peaks are identified with a peak-finding algorithm that enforces a minimum inter-peak separation of 10 bins and rejects peaks below 5 % of the global histogram maximum. These parameters were found empirically, by looking at the histogram and then enforcing that peaks that are only one bin apart ought to be counted as the same peak. The position uncertainty of each peak is estimated as half its full width at half maximum to account for this previous bunching, providing some weights for the calibration fit.

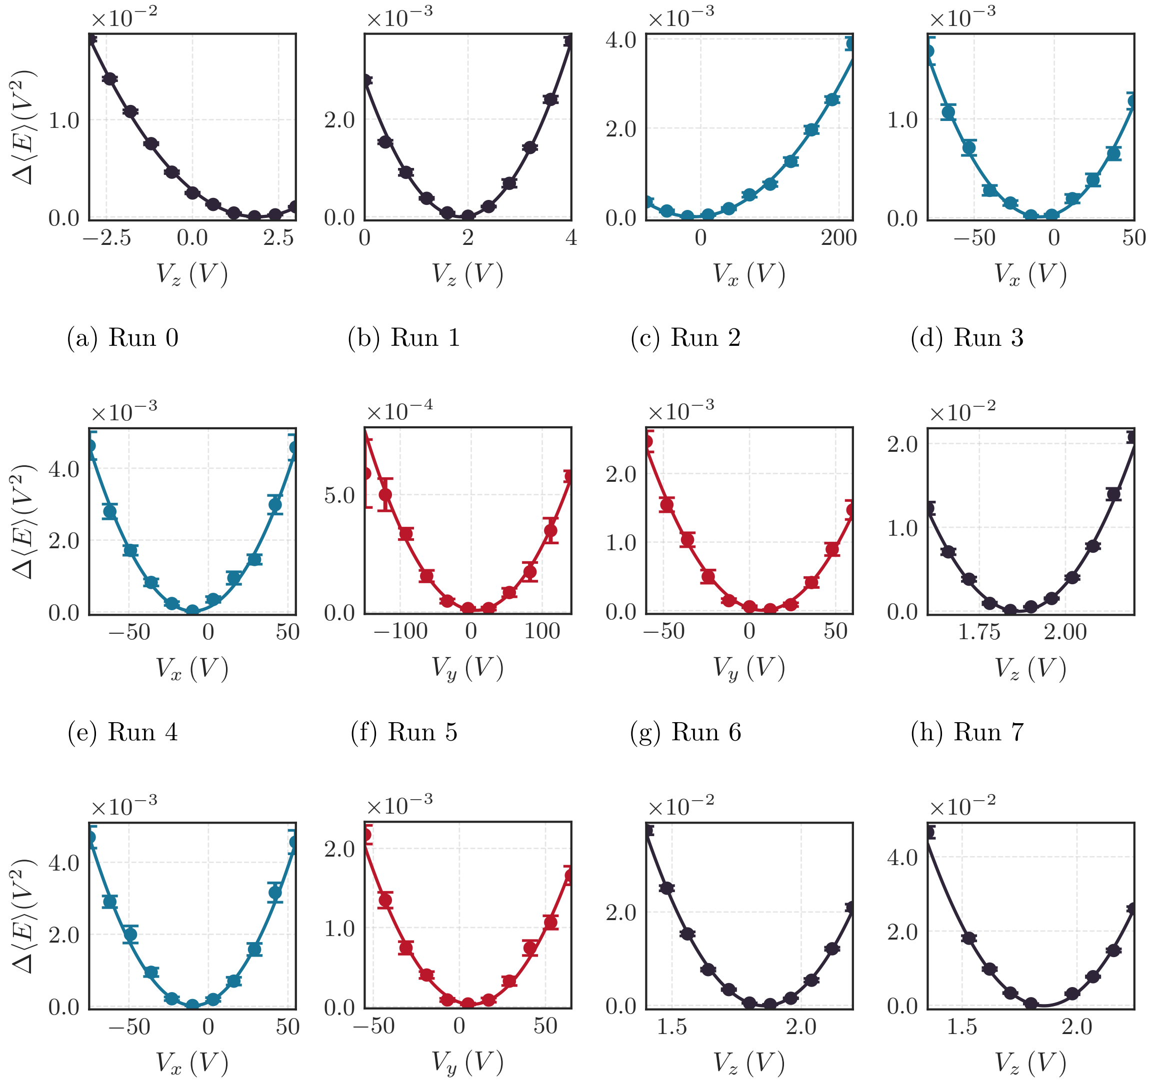

(a) Run 0 (b) Run 1 (c) Run 2 (d) Run 3

(e) Run 4 (f) Run 5 (g) Run 6 (h) Run 7

(i) Run 8 (j) Run 15 (k) Run 22 (l) Run 23

Figure 4.12.: **Representative parabolic fit plots for the voltage scans in Table 4.3.** Note that runs 8, 15, and 23 are shown as the representative plots for their respective averaged ensembles. The analysis is done via mode analysis, where the increase in mean energy of the respective mode is plotted against the compensation voltage. The solid line is a fit through the data using the function $a(V - V_{\text{opt}})^2 + b$, where a, V_{opt} and b are fitting parameters. The error bars show the standard error of the mean of five drops at each voltage.

Because the particle occasionally jumps by more than one elementary charge at once, the histogram peaks are not necessarily uniformly spaced. To assign correct integer charge-state indices, an unweighted linear fit to a contiguous reference subset of peaks is first used to estimate the mean voltage step per elementary charge δV. Consecutive peak separations are then divided by δV and rounded to the nearest integer, yielding the number of charges between successive peaks; absolute charge-state indices are constructed by accumulating these integer steps. A final weighted least-squares fit of peak voltage against charge-state index, using the FWHM-derived uncertainties as weights, then yields the calibration slope m and offset c, with formal uncertainties from the covariance matrix of the fit. Only peaks with charge-state indices in the range $[1, 15]$ are included in this fit, as levels outside this window are statistically less well sampled and not as clearly visible. The resulting calibration and fit are shown in Figure 4.14.

It is therefore important to run the calibration in a way that every charge level visited is well-resolved in the histogram. Setting a low trip-current limit of 1 mA and manually resetting the power supply's resettable current fuse after each charge change allows the particle to remain at a set charge level for a longer time, while not changing the overall calibration result. The current limit also is observed to be roughly proportional to the jump height distribution. For higher currents, the bigger jumps seem to appear more frequently. An alternative approach is simply to wait for the particle to settle on all levels more often and resetting the current fuse automatically. Since the analysis relies only on the *positions* of the histogram peaks and not on their heights, the two approaches yield equivalent calibration results.

The calibration is valid as long as both the drive amplitude and detection remain unchanged, and can be transferred to different particles provided these conditions are met. In this setup we find a typical charge-to-voltage ratio of $(257 \pm 2)\,\mu\text{V/e}$ with an offset of $(10 \pm 20)\,\mu\text{V}$ for the axial direction, which allows us to determine the absolute particle charge by dividing the measured lock-in amplitude (minus the offset) by the slope and rounding to the nearest integer. From this we find that natively after loading, the particles usually have on the order of 20 e charge.

Turning back to the histogram, we can now use the slope to convert all lock-in levels to charge states, which then allows us to see how long the particle spends in each charge state during the calibration process, as shown in Figure 4.15. This can be used to determine the validity of the calibration, as the charge levels from the fitted slope roughly match the peaks in the histogram.

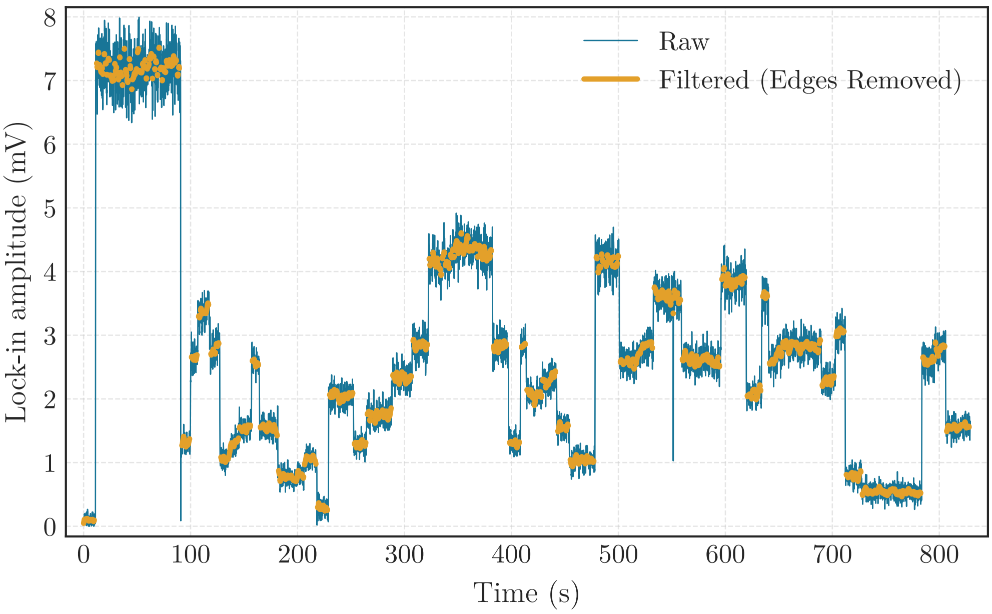

Figure 4.13.: **Time trace of the charge calibration process**, showing the lock-in level changes as the particle charge state is modified. The overlay is the data that is used in the further analysis. Using differences between the datapoints, we can detect the edges and remove the in-between points to create a clean signal of the charge states that were visited during the calibration process.

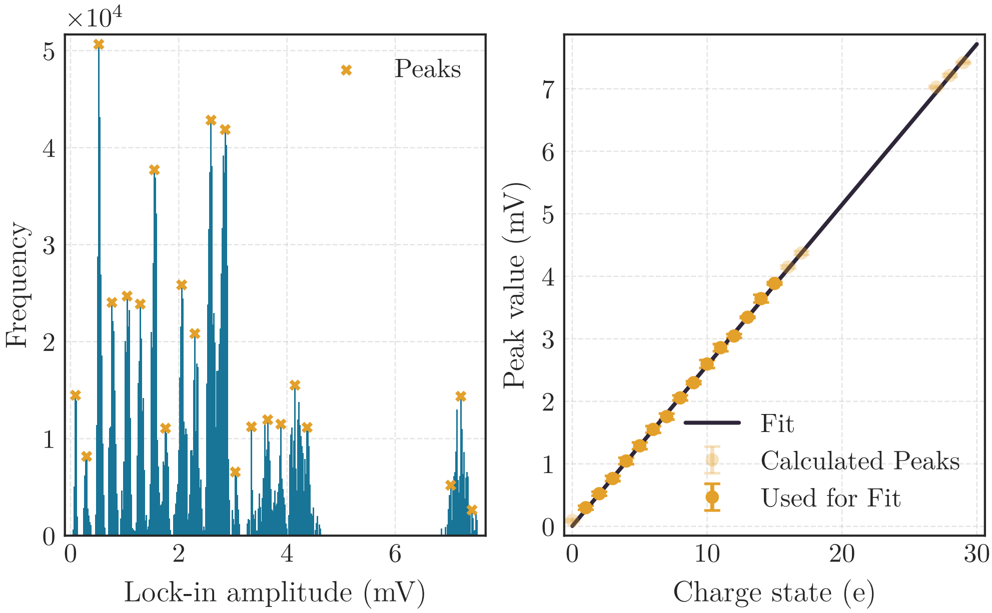

Figure 4.14.: **Charge to lock-in voltage ratio calibration. Left:** A simple histogram of Figure 4.13 with found peaks indicated by a cross. **Right:** The same histogram's peaks fitted to a linear function. By fitting the histogram of the charge states obtained during the calibration process, we can extract the charge to voltage ratio as $(257 \pm 2)\,\mu\mathrm{V/e}$ with an offset of $(10 \pm 20)\,\mu\mathrm{V/e}$. The different peaks correspond to different charge states of the particle, with the lowest level being uncharged and then increasing by integer steps of elementary charges. The fit is done using a linear function, as the lock-in level changes linearly with the charge state.

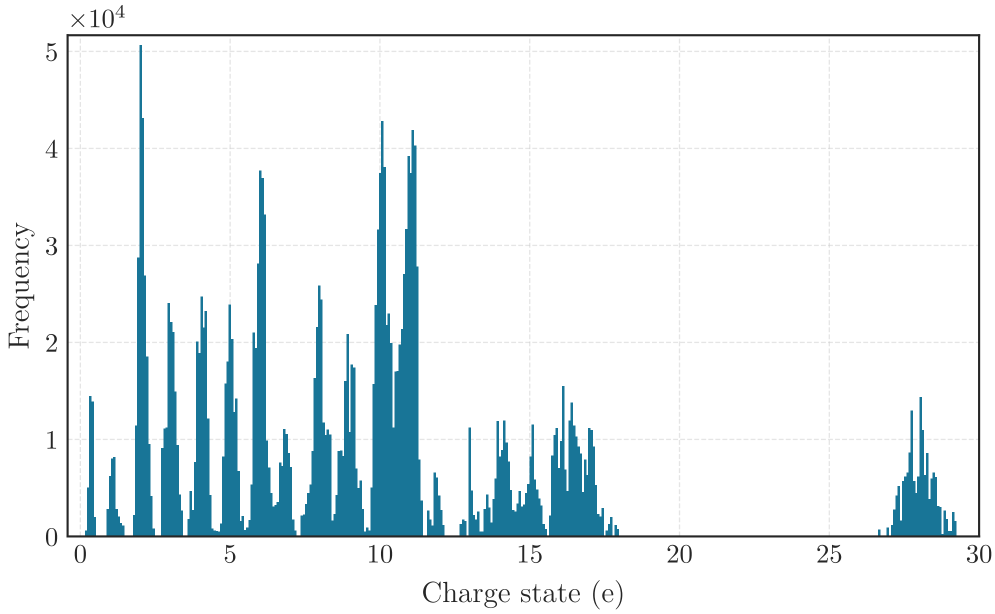

Figure 4.15.: **Calibrated histogram of the charge states obtained during the calibration process.** Cross checking the calibration result by comparing the measured charge states with the position of the peaks in the histogram. The peaks in the histogram should roughly match the charge states obtained from the fitted slope, which is the case here, indicating a valid calibration.

4.7.2. Environment charge manipulation via low energy electrons

Utilizing a different power supply, we can also deposit some very low energy electrons into the vacuum chamber by thermionic emission from a hot tungsten filament at low pressures (1×10^{-7} mbar), much like a light bulb and shown in Figure 4.16. Setting a bias voltage on the electrodes then would allow to direct these electrons towards or away from the nanoparticle trapping region and surfaces in the vacuum chamber, thereby changing the environment charge.
By monitoring the compensation voltages needed to minimize the free evolution energy increase as in the compensation voltage scans, we can see how the environment charge changes. This method could be particularly useful when the compensation voltage scans do not show a clear minimum, as it allows to shift the minimum into the scan range. The particle charge state remains constant during this process. The method is illustrated in Figure 4.17, where the compensation voltage in z is monitored over time while the filament is turned on and off at different voltages on the electrodes.

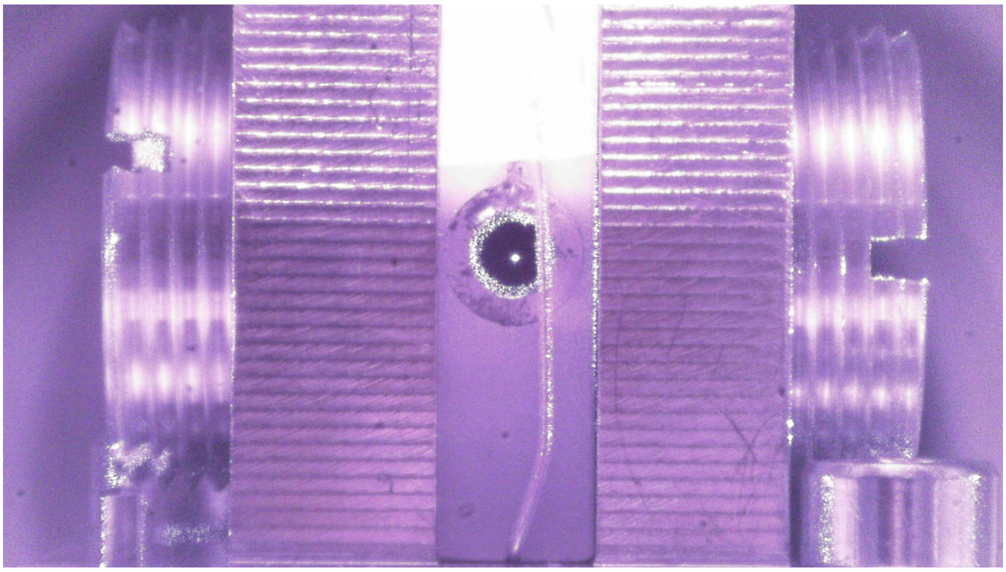

Figure 4.16.: **Environment charge manipulation via thermionic emission from a hot filament.** This image shows the glow from the filament in the vacuum chamber captured with a camera through the side window.

The effect of moving the compensation voltage minimum is clearly visible, indicating a change in environment charge, but it is important to note that the environmental charge state cannot be as easily controlled for the radial directions. While the process, which we have convinced ourselves is the reason for the change in the V_{opt} in z, should work for x and y, we were not able to show this conclusively. We attribute this to the geometry of the electrodes and the holder. As the filament emits electrons which then need to be guided to the electrodes, it can make some directions more favourable than others. Additionally complicated is that a shifted environment charge changes the minimum compensation voltage in z, thereby changing the bias for the x and y scans as well, which then in turn leads to different ranges that can be scanned in those two directions. It is therefore hard to quantify the effect in x and y as well, but it still should be possible to shift the minimum in those directions as well. We did not explore this in detail for this thesis as there is a bigger caveat to this method.

Charging the environment in this way can lead to long term drifts of the environment charge state, as the low energy electrons can get stuck on surfaces inside the vacuum chamber and then slowly discharge over time. By monitoring the optimal compensation voltages over long times, one can see drifts like those of an exponential decay as one would expect from a charged capacitor discharging over a resistor. This effect can be mitigated by only using the filament for short times to shift the compensation voltage minimum into the scan range, then waiting for the environment charge to stabilise enough for the required measurement accuracy. In Figure 4.18, the stability of the environment charge after a flash of the filament is shown. It is also important that the charge state of the environment is stable for long times, as this would void any compensation voltage scan, which rely on the environment charge being constant to be valid after scanning, since they are intended to be preparatory in advance of a measurement that can then take a long time to perform. A drift in compensation voltage would therefore be inversely proportional to maximally permissible measurement time, as the longer the measurement, the more drift can happen during the measurement, leading to a larger error.

To analyse this change in charge, we utilise the same method as for the compensation voltage scans, performing scans multiple times over a long period and monitoring the change in compensation voltage minimum, tracking the range such that the minimum is always visible in the scan. By doing this, we can see how the environment charge changes over time after being manipulated by the filament. The resulting data is shown in Figure 4.18. Then an exponential fit is calculated for the scan batches taken after the flash to extract the time constant of the decay of this charge, which is found to be around 4 h. A linear fit is also calculated for the first batch of scans taken before the flash. Both scans have their confidence intervals calculated and plotted as well, showing that the environment charge is stable over hours without manipulation, but then drifts over time after being manipulated.

Figure 4.18 is clearly indicating that the environment charge is stable over hours if not manipulated at all and discharges slowly over time after being manipulated and moved to a different charge state. Since it is an exponential decay and seemingly does not settle to the original value, this may not be a complete dealbreaker, but it does mean that one would have to wait sufficiently long after manipulating the environment charge to ensure that the compensation voltages are stable enough for the required measurement accuracy.

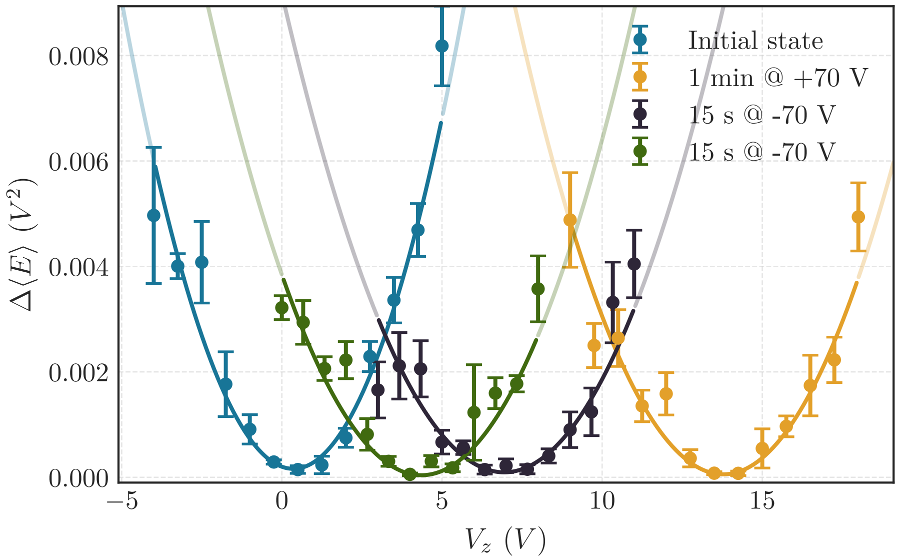

Figure 4.17.: **Environment charge manipulation via thermionic emission from a hot filament.** By turning on the filament at different bias voltages of the feedback electrodes, we can see the compensation voltage in z change over runs, indicating a change in environment charge due to the thermionic emission of electrons from the glowing filament. The Bias voltage held at the electrode during the flash is given in the legend with the corresponding time the flash was running. The particle charge remains constant during this process. For each run, the particle is released for a free evolution of 10 µs at 13 compensation voltages spanning the scan range. At each voltage, the detector signal is divided into 5 bins and a power spectral density (PSD) is computed. Integrating the PSD in a window around the z-mode frequency (86 kHz) yields the motional energy increase $\Delta\langle E\rangle$ accumulated during the free evolution. The procedure is conceptually the same as for the compensation voltage scans in Section 4.6 or also the stability scans in Figure 4.18.

4.7.3. Aside: Discharging the Environment

For this experimental setup multiple ways to discharge the environment were tested. Ideally, the previous methods and operations might leave the environment charged up, but then there would be a way to discharge it again. A few methods already are well established in literature, like irradiating the environment with UV light [57, 58], but also flowing ionised nitrogen gas through the chamber [59, 60] is also used for LIGO test masses [61], however both methods did not work well in this setup.
The UV light method was attempted with a 300 nm UV lamp placed close to the chamber on the outside, shining it through the trapping laser window, but no significant change

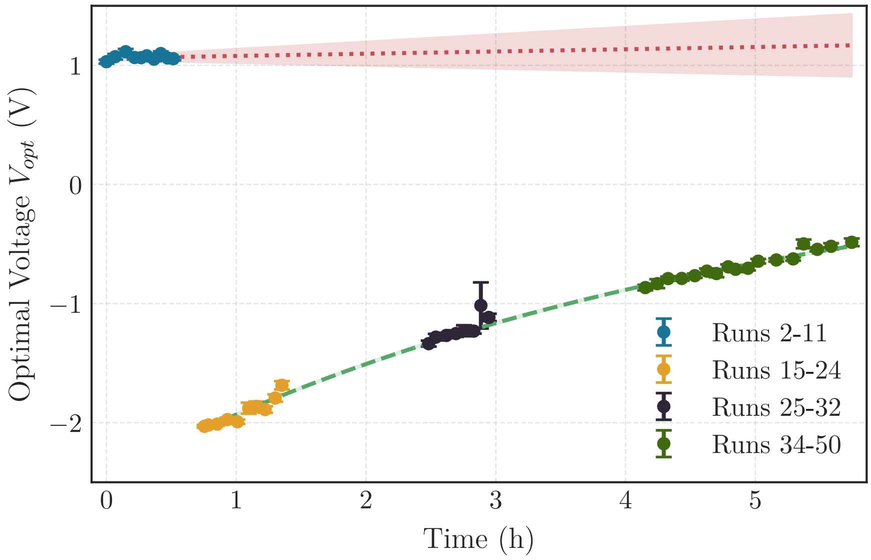

Figure 4.18.: **Stability of the environment charge after a flash of the filament.** The compensation voltage in z is monitored over time, showing the relaxation of the environment charge after manipulating it by low energy flashing. It can be seen that the environment charge is 1) stable over hours if not manipulated at all (Runs 2-11, blue) and 2) Discharges slowly over time after being manipulated, as indicated by the linear fit through the data (red dotted line). The decay can be fitted to an exponential, just like a capacitor discharging over a resistor. From this fit a time constant *RC* of 4 h is extracted (Runs 15-50, excl. 33, green dashed line). The shaded areas indicate a 1σ confidence interval, showing that the environment charge is stable over hours without manipulation. The error bars on each point are the errors of the parabola fit.

in compensation voltages could be observed after extended illumination times of multiple minutes. This could be due to the geometry of the chamber and the placement of the lamp, as well as the materials used inside the chamber as well as the glass, which might not be very transmitting at this wavelength.
Also any method that would lead to discharged environment at high pressure would then have its effect masked by triboelectric charging when the chamber is being evacuated again, as the gas flow through the vacuum pumps and tubing would then lead to static charge buildup on surfaces inside the chamber again.
The only somewhat reliable method to discharge the environment was to let the chamber sit for long times (hours) at high pressures (10^{-1} mbar) and hope that the environment slowly discharges itself through natural processes like gas collision. This method is not very practical, as the environment charge state is not reset after some long charge calibration scans but remains at some offset level for short times left at high pressure. After time however, the compensation voltages do tend to drift back towards zero, indicating a slow discharge of the environment. During this time, depositions in the chamber also occur, making the subsequent pumpdown to high vacuum take longer. When pumping down, also here we have triboelectric charging from the gas flow, which can lead to a different environment charge state than before,but this method can work to come back from a state in which the environment charge was un-workably high, such as after the charge calibration measurements.

4.8. Optimal compensation voltages as a function of free evolution times

As shown in Equation (2.23) the mean oscillation energy scales as $F_\mathrm{z}^2\tau^4$ for $\tau \gg \Omega^{-1}$ in the presence of a static force F_z. To measure this dependence experimentally, we performed compensation voltage scans for different free evolution times τ and extracted the optimal compensation voltages that minimize the mean energy increase after free evolution by fitting a parabola $a(V - V_\mathrm{opt})^2 + b$ to each scan (as we also did before in Sections 4.3, 4.3.1 and 4.6). The force is proportional to the compensation voltage, so we see the τ^4 scaling from the scale factor a of the parabola.
This dataset gives insight into the variation of the force, as well as significant kicks from the execution of the protocol to the particle. By plotting the optimal compensation voltage as a function of τ we expect to see a flat line if there are no fluctuations of the force and no significant kicks.
The data agrees well with a constant value, as shown in Figure 4.19, indicating that the average force acting on the particle does not change significantly over free evolution times. Fluctuating forces would show up here and lead to a trend in the optimal voltage, as for each drop duration the voltage would be different. Since we also cannot see asymptotic behaviour for long times, we can conclude that there are no significant kicks to the particle from shutting off the trap and the cooling during the protocol. These would manifest mostly at short times, as the ratio between a force from such a kick and the static force would be highest.

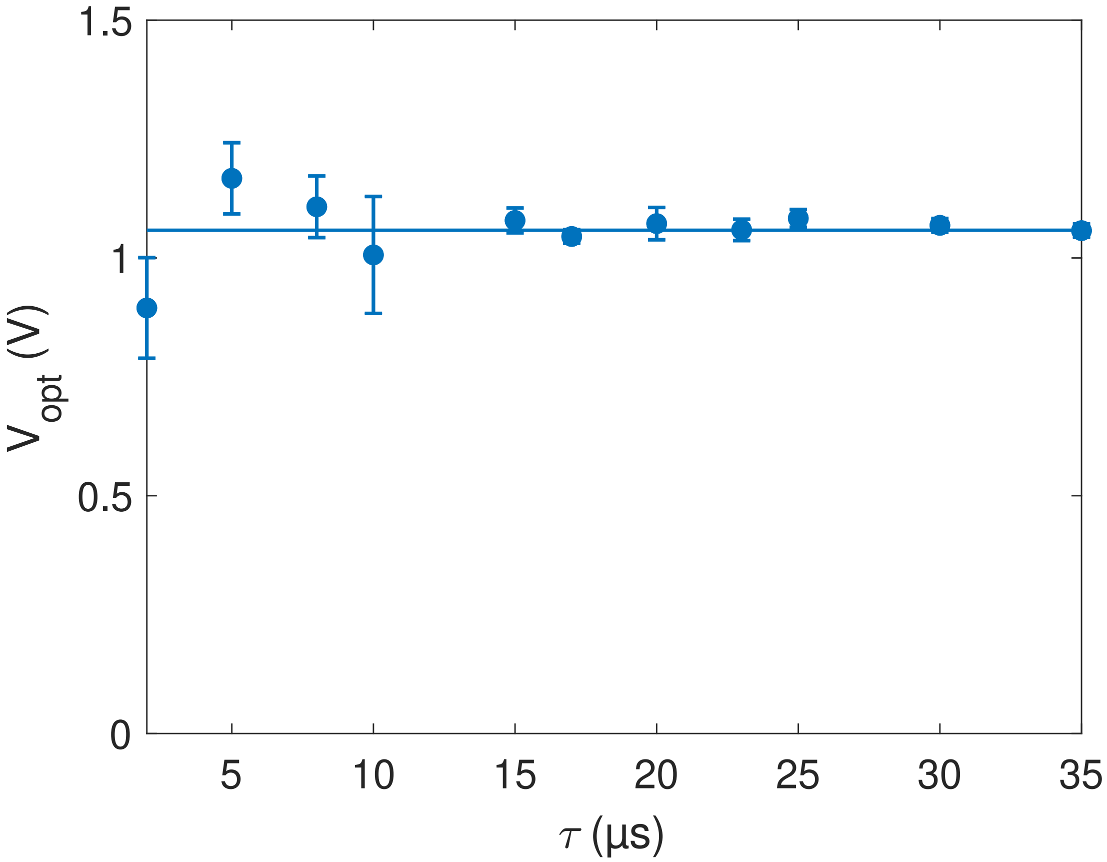

Figure 4.19.: **Minimum compensation voltages as a function of free evolution times.** Plotted here is the optimal compensation voltage V_{opt} obtained from the parabola fits to the compensation voltage scans for different free evolution times τ. The data shows good agreement with a constant value, indicating that the average force acting on the particle does not change significantly over time and that there are no significant kicks to the particle from executing the protocol. The error bars are the reported standard errors of the fit parameters. Figure credit: Y. Y. Fein

4.9. Mean energy increase as a function of free evolution time

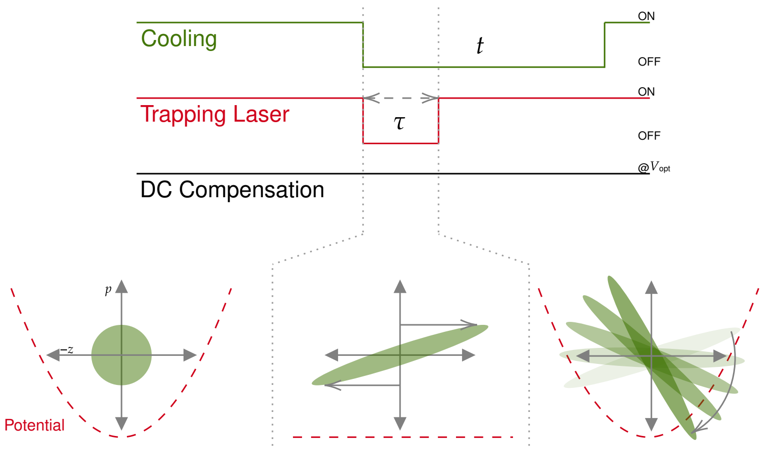

Figure 4.20.: **Pulse for the free evolution time scan protocol.** The particle is first cooled via to a low energy state. Then, the trap and feedback are switched off simultaneously for a variable free evolution time τ, after which the trap is re-enabled and the cooling remains off for the remaining duration t. The particle's motion is then recorded after recapture to give enough data to extract the mean energy after the free evolution. During the free evolution, the phase space is sheared and after recapture, the state rotates around the origin due to the harmonic potential. To prevent displacement, the DC-Biases are held at their respective optimum values. This procedure is repeated for different free evolution times τ to extract the mean energy increase as a function of τ.

We can quantify how well DC forces are compensated by measuring the mean energy as a function of τ. In the presence of a constant residual force F, the full expression for the mean energy ratio after a free evolution of duration τ is (copied from Equation (2.24) for convenience):

$$\frac{\langle E(\tau)\rangle}{\langle E_0\rangle} = 1 + \frac{\Omega^2\tau^2}{2} + \frac{F^2\tau^2}{2mE_0} + \frac{F^2\Omega^2\tau^4}{8mE_0}. \tag{4.4}$$

Grouping the two τ^2 contributions gives $1 + \left[\Omega^2 + F^2/(mE_0)\right]\tau^2/2$, where the harmonic term Ω^2 exceeds the force term $F^2/(mE_0)$ by several orders of magnitude for the residual forces present here. The dominant scaling is therefore $1+\Omega^2\tau^2/2$. We can see the residual force best as the quartic term $F^2\Omega^2\tau^4/(8mE_0)$, which only becomes distinguishable from the pure harmonic prediction at longer drop times. Both curves are shown in Figure 4.21. For this measurement, the 3D compensation voltages were applied.
To measure the free-evolution energy gain as a function of free evolution time, each acquisition file contains the particle's homodyne position signal, digitised at $f_s = 25,\text{MHz}$, with the trigger located at the sample centre. From each file, here too, pre- and post-

evolution windows are extracted. The post-evolution segment $x_2(t)$ begins a fixed dead-time of 3 µs after recapture. This segmentation is performed for all 150 statistically independent runs at each of the 18 drop times $\tau_i \in 0, 1, 2, 4, \ldots, 100$µs, resulting in 2700 pre-drop and post-drop segments for analysis.

To extract the oscillation amplitude in the presence of multiple mechanical modes, the post-drop segment is fitted to a triple-sine model with a linear drift term,

$$x_2(t) = \sum_{k=1}^{3} A_k \sin(\omega_k t + \varphi_k) + c_0 + c_1 t, \tag{4.5}$$

capturing the axial mode near 91 kHz and the two radial modes near 271 kHz and 305 kHz, while the linear term absorbs slow DC drift unrelated to the mechanical oscillation. The three frequency guesses are extracted individually for each run from periodogram of a longer window. Once the post-evolution frequencies are determined, the pre-evolution segment is fitted with a single-sine model at the now determined axial frequency ω_1, giving us the pre-evolution amplitude A_1' for that mode. The detail of this fitting is described in Section A.2.5.

Rather than using the raw signal variance, which would include contributions from all three mechanical modes and the slope, the mode variance relevant for the energy comparison is reconstructed analytically from the fit amplitudes as $\sigma_k^2 = A_k^2/2$, corresponding to the time-averaged mean-square displacement of a pure sinusoidal oscillation. The relative mean energy after a drop of duration τ_i is then computed from the ensemble of 150 runs as

$$\frac{\langle E \rangle}{\langle E_0 \rangle}\bigg|\tau_i = \frac{\overline{\sigma_2^2}(\tau_i)}{\overline{\sigma_1^2}(\tau_i)}, \tag{4.6}$$

where the overbar denotes the ensemble mean after masking runs for which the fit did not converge. The error is quantified by non-parametric bootstrap resampling (Nboot $= 300$, see Section A.2.8), giving a 1σ confidence interval on the ratio.

The resulting energy ratios are shown in Figure 4.21 as a function of drop time, and a theoretical curve for the expected energy increase with a force of 1.9×10^{-18} N is plotted as well.

The data shows good agreement with the expected evolution in the absence of an external force, demonstrating the efficacy of DC compensation. The mean displacement using this compensation method along z is confined to less than 10 nm for free evolutions up to 100 µs.

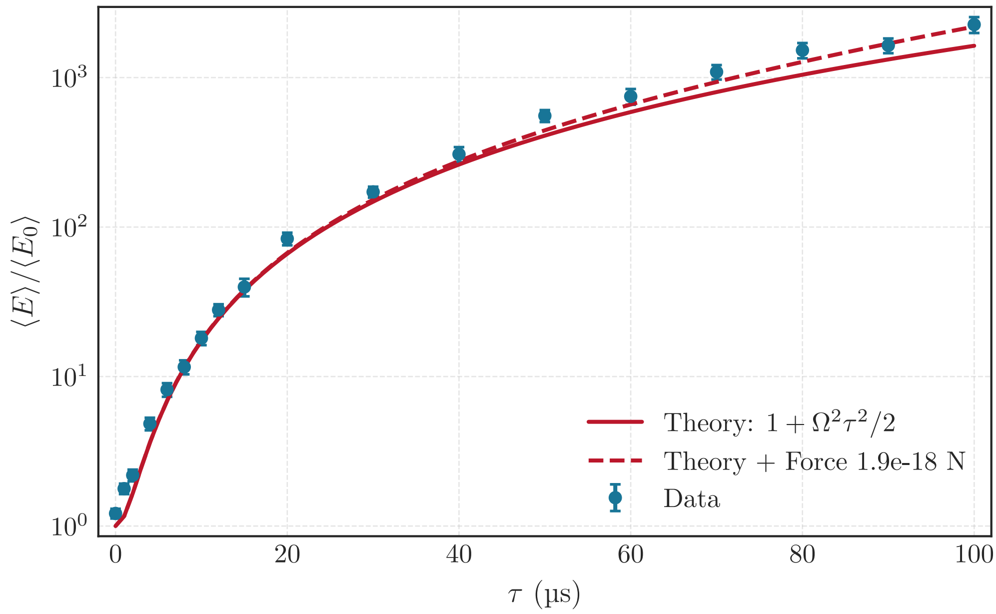

Figure 4.21.: **Mean energy increase as a function of free evolution time at low pressures (** 8×10^{-8} mbar **)**. From this data we see that the mean energy increase scales mostly as $1+\Omega^2\tau^2/2$ and is consistent with a residual force of 1.9×10^{-18} N. The mean energy is extracted by fitting the squared amplitudes of a sine fit where each point is 150 repetitions. The error bars show the standard error of the mean of all fit amplitudes. The solid yellow line shows the expected evolution of the mean energy in the absence of an external force, showing good agreement with the data.

4.10. Position standard deviation growth as a function of free evolution time

This measurement uses the same pulse sequence and dataset as the preceding mean energy analysis (Figures 4.20 and 4.21).

The mean energy ratio $\langle E\rangle/\langle E_0\rangle$ characterises how much energy the particle gained on average, but reduces the full phase-space distribution into a single scalar value (Mean energy ≈ mean distance from phase space origin). As discussed in Section 2.4.1, the direct figure of merit is the actual spatial extent of the motional state: the position standard deviation $\sigma_z(\tau)$ of the ensemble after a free evolution of duration τ. This is the quantity that directly relates to the spatial delocalisation of the particle.
The spread arises from the same phase randomisation that also underlies the mean energy calculation. Each repetition is recaptured with a phase φ_j drawn from a distribution set by the initial thermal state, so the 150 trajectories form a cloud of points in phase space. While the mean energy depends only on the oscillation amplitudes A_j, the position standard deviation depends also on the phases, and in particular on how the cloud is oriented relative to the position quadrature at any given moment after recapture. Because the cloud rotates in the harmonic potential, its projection onto the position axis oscillates, and its maximum over one oscillation period defines the maximum state size.
The axial trajectory for each run j is reconstructed from the triple-sine fit parameters $(A_j, \omega_j, \varphi_j)$ on a dense time grid of 1000 equally spaced points over the post-drop window $[0, t_{\mathrm{post}}]$. Computing the ensemble variance across the 150 runs at each time point,

$$\mathrm{Var}_{\mathrm{ens}}(t;\, \tau) = \mathrm{Var}_j[A_j \sin(\omega_j\, t + \varphi_j)]\,, \tag{4.7}$$

yields a time-dependent quantity that oscillates as the phase-space cloud rotates in the harmonic potential. Its maximum,

$$\sigma^2_{\max}(\tau) = \max_t\ \mathrm{Var}_{\mathrm{ens}}(t;\, \tau), \tag{4.8}$$

captures the characteristic state size for that drop time regardless of the shearing angle θ. Dividing by the detector calibration constant α (see Section A.1) gives the position standard deviation

$$\sigma_z(\tau) = \frac{\sqrt{\sigma^2_{\max}(\tau)}}{\alpha}. \tag{4.9}$$

Error bars are obtained from bootstrap resampling ($N_{\mathrm{boot}} = 300$, see Section A.2.8) using the same procedure as the mean energy analysis, giving a $1\,\sigma$ confidence interval on σ_z. Two theoretical curves are shown in Figure 4.22. The theoretical curve (dashed) propagates the measured mean pre-drop variance $\bar{\sigma}_0^2$ through the free covariance-matrix evolution of a thermal state,

$$\sigma_z^{\mathrm{th}}(\tau) = \frac{\sqrt{\bar{\sigma}_0^2\,(1 + \Omega^2\tau^2)}}{\alpha}, \tag{4.10}$$

which gives the maximum position-quadrature extent of the sheared phase-space ellipse as a function of τ. At long drop times this simplifies to the asymptotic expression (dotted).

$$\sigma_z^{\mathrm{asymp}}(\tau) \approx \frac{\sqrt{\bar{\sigma}_0^2}\,\Omega\tau}{\alpha} = z_{\mathrm{zpf}}\sqrt{2\bar{n}+1}\,\Omega\tau, \tag{4.11}$$

where $z_{\mathrm{zpf}} = \sqrt{\hbar/(2m\Omega)}$ is the zero-point fluctuation amplitude and $\bar{n}$ is the mean phonon number of the initial thermal state.

The measured state size, shown in Figure 4.22, follows the theoretical curve (Equation (4.10)) closely across the full range of drop times, confirming that the state expands as expected for a freely evolving thermal ensemble. At short τ the data lies above the asymptotic line (Equation (4.11)), which is expected: the asymptote is the large-τ limit of Equation (4.10), obtained by dropping the $+1$ inside the square root ($\sqrt{1+\Omega^2\tau^2} \approx \Omega\tau$). This approximation fails when $\Omega\tau < 1$ because the initial thermal size σ_0 still contributes comparably to the final size. This becomes especially clear when calculating the position standard deviation for both Equations (4.10) and (4.11) at $\tau = 0$. We can see the position delocalisation growing to 4 nm from free evolutions of 100 µs.

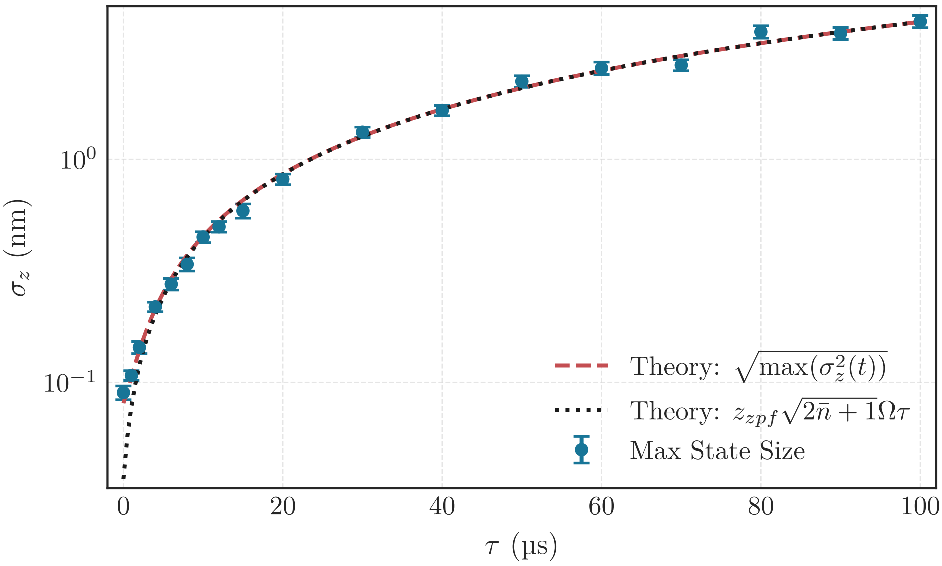

Figure 4.22.: **Position standard deviation growth as a function of free evolution time.** Same dataset as Figure 4.21, reanalysed to extract the position standard deviation $\sigma_z(\tau)$ rather than the mean energy. Each data point is the maximum ensemble spread over one oscillation period, extracted from 150 sine-fitted trajectories and converted to nanometres via the detector calibration. The dashed line is the theoretical prediction $\sigma_z^{\text{th}} = \sqrt{\bar{\sigma}_0^2(1+\Omega^2\tau^2)/\alpha}$ using the measured pre-drop variance as input; the dotted line is its large-τ asymptote $\sigma_z^{\text{asymp}} \approx \sqrt{\bar{\sigma}_0^2\,\Omega\tau/\alpha}$. The data follows the theoretical curve closely across the full range. The deviation from the asymptote at short τ is expected because the initial state size σ_0 remains comparable to the growth term when $\Omega\tau < 1$. The delocalisation reaches 4 nm at $\tau = 100$ µs. Each point is 150 repetitions; error bars are $1\,\sigma$ bootstrap confidence intervals (see Section A.2.8).

4.11. Long Free Evolutions

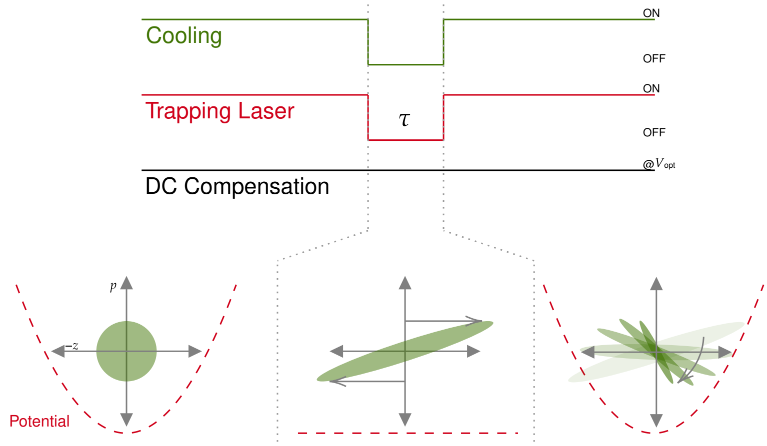

Figure 4.23.: **Pulse sequence for the long free evolution protocol.** The particle is first cooled via to a low energy state. Then, the trap and feedback are switched off simultaneously for a long free evolution time τ, after which they are re-enabled. During the free evolution, the phase space is sheared and after recapture, the state rotates around the origin due to the harmonic potential. To prevent displacement, the DC-Biases are held at their respective optimum values. This procedure is repeated for ever-increasing free evolution times τ, usually until the particle can no longer be recaptured. The measurements are all single shot as τ is too large to reliably repeat the measurement often, and the goal is finding the new limit instead of reliably performing the free evolutions in rapid succession.

Although we set out to push the limits of free evolution times, we have not yet shown free evolution times *far beyond* the previous state of the art and not in the order of magnitude needed for the one protocol introduced in Section 1.2. To reiterate the point already made multiple times throughout, the longer the free evolution can last, the better, as it is inversely related to the magnitude of residual static forces. Additionally, as mentioned when discussing the figures of merit in Section 2.4, pushing the boundaries of free evolution and looking at the spectrograms of the motion after recapture is a good way to find paths to further improve the free evolution times. The pulse sequence for the long free evolution trials is shown in Figure 4.23, and differs from the previous measurements mostly in the fact that the cooling is re-enabled after recapture, as the motion should be damped down from the highly excited state immediately after recapture.
The main way to analyse these long free evolutions is to look at the spectrogram of the motion and also of the cooling signal after recapture. While this is not a very quantitative approach, it allows to see the excitation of the motional modes and their cooling back down after recapture, as well as any frequency shifts that might occur due to the particle exploring nonlinearities from Equation (2.21) and Figure 4.8 in the trap.

The measurement was performed using the bandpass cooling method from Section 3.1.6. To calculate the spectrogram, we record four channels on the oscilloscope, called Channel A-D. We want to know both the motion of the particle and also the cooling signal produced after filtering this motion that is being fed back to the particle for cooling. Channel A is the BW homodyne detection of the particle motion, set to a range of 5 V, Channel B and C are the x/y FW detections respectively, set to a range of 2 V, and Channel D is the z cooling signal, which is the output of the cooling filter, set to a range of 2 V. The sampling rate is set to 20 MHz but actually achieves 10.4 MHz. In the processing, we take a symmetric time window around the trigger point, which we zero our time axis on, and then calculate the spectrogram using the `scipy` package with Hann windows the size of `nperseg=16384` and an overlap of 90 %. The `nperseg` is the window size that the spectrogram is calculated on, so larger window sizes give better frequency resolution, at the cost of time resolution. The overlap is the amount of overlap between the windows, making the spectrogram smoother, at the cost of computation time. The color bar on the lowest pixels is a debugging tool to see if the voltages were out of range during the measurement, where green indicates that the voltages were within range, while red indicates that they were out of range.
In Figure 4.24 we see such a spectrogram after some long free evolutions of 500 µs and 1000 µs, showing the particle's motion after recapture.
The excitation in the motional modes $\Omega_{x,y,z}/(2\pi) \approx [300, 260, 90]$ kHz is clearly visible being cooled down as the feedback is turned back on after 6 ms, but the particle remains stably trapped after both of these evolution times, indicating good compensation of DC forces and surpassing the previous state of the art mentioned in Section 1.3.3.
For the 1000 µs free evolution, we can see some harmonics of the motional modes, also indicating one of the problems the cooling faces for strongly excited motion, as these harmonics also leak into the cooling signal as can be seen in (h) of Figure 4.24.

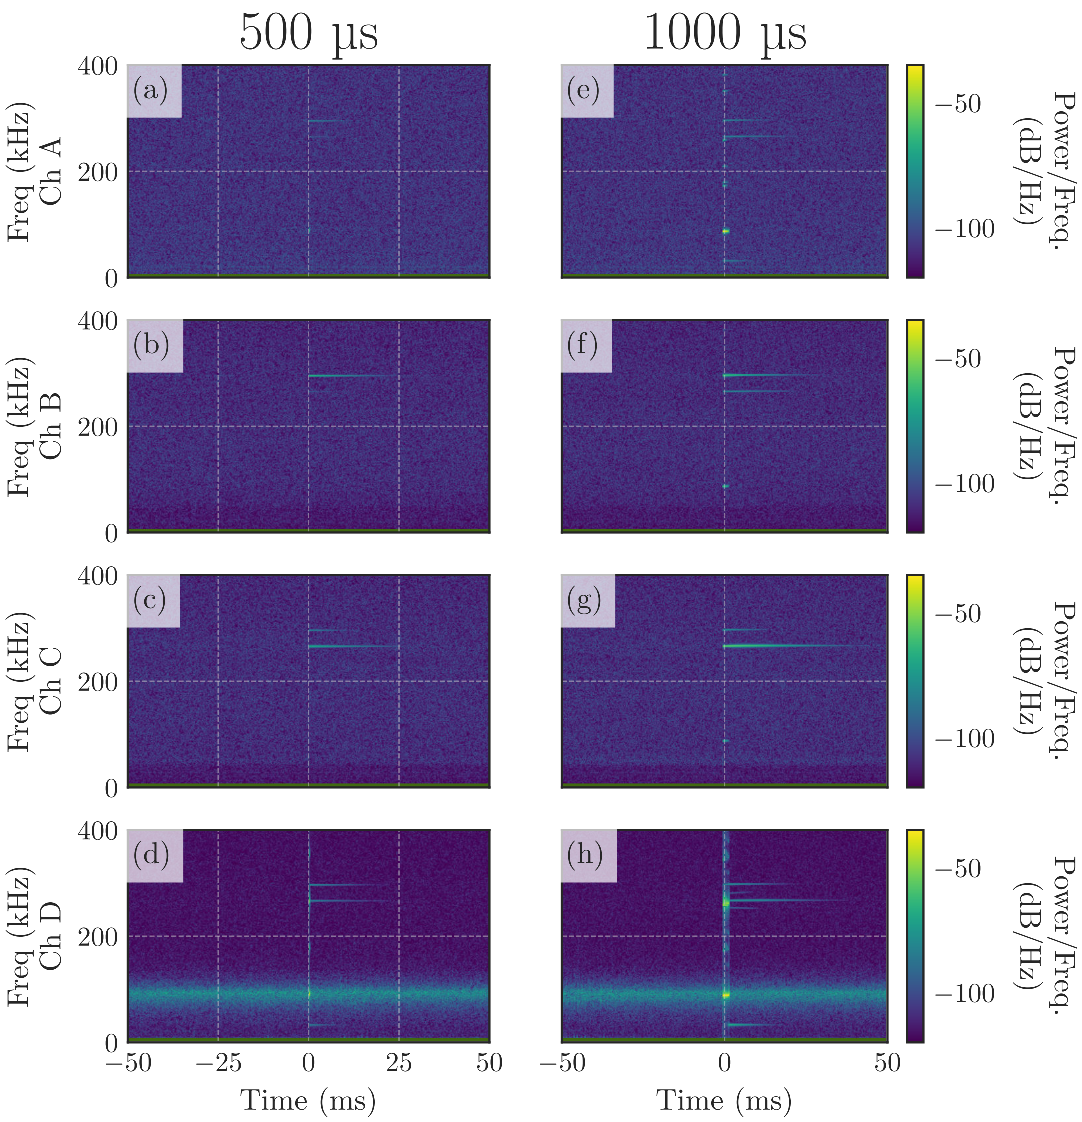

Figure 4.24.: **Spectrograms of a** 500 µs, 1000 µs **free evolution followed by recapture.** The trap and feedback are switched off at $t = 0$ and re-enabled after the free expansion time plus. Excitation and re-cooling of the motional modes $\Omega_{x,y,z}/(2\pi) \approx [300, 260, 90]$ kHz is visible upon recapture. Signals are recorded with a sampling rate of 10.4 MHz. **(a),(e)**: BW Homodyne signal in 5 V range, **(b),(f)**: x FW detection, **(c),(g)**: y FW detection, **(d),(h)**: z cooling signal all in 2 V range. The faint colored line at the very bottom of the spectrogram is an indicator if the voltages were out of range during the measurement, where green indicates that the voltages were within range, while red indicates that they were out of range.

4.12. Squeezing as a function of realignment pulse time

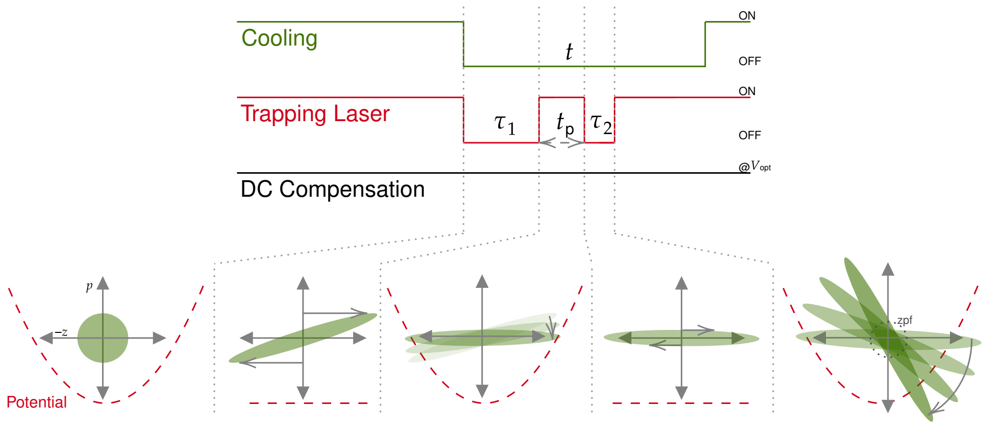

Figure 4.25.: **COM-Mode Squeezing as a function of realignment pulse time.** The particle is first cooled to a low energy state. Then, the trap and feedback are switched off simultaneously for τ_1. the state expands and shears in phase space. After τ_1, a brief trap pulse of duration t_p is applied, which imparts a rotation in phase space that aligns the squeezed quadrature with the position quadrature at the moment of the second release, shearing now perpendicular to the minor axis of the ellipse. The particle then evolves again for τ_2 before recapture. By varying the pulse duration t_p, the state can be over (under) rotated for times longer (shorter) than the optimum.By comparing how a ground state would have expanded, we can dertermine if our state was indeed squeezed below z_{zpf}

Having now shown that long free evolutions are possible with the methods developed, we can also explore pulsed protocols that show the quantum nature of the system, such as squeezing. As laid out in Section 2.2, squeezing is a reduction of the variance of one quadrature below the standard quantum limit, at the cost of an increased variance in the conjugate quadrature. This effect cannot be explained by classical physics and is a unique property of quantum mechanics. This is especially interesting as our experiment is not in the ground state, so a squeezing below the ground state variance would not only be a demonstration of the quantum nature of the system, but also a demonstration of efficacy of our free evolution for squeezing in general.

The squeezing notebook implements an asymmetric protocol: the particle first evolves in free evolution for $\tau_1 = 20\,\mu s$, receives a brief trap pulse of duration t_p, then evolves again for $\tau_2 = 15\,\mu s$ before recapture. The pulse duration now is sub-period ($t_p \ll T/2 \approx 5.6\,\mu s$). Its role is now to impart a rotation in phase space that aligns the squeezed quadrature with the momentum quadrature at the moment of the second release, so that the squeezing can be observed in the position quadrature after the second free evolution via time-of-flight measurement.

The analysis of the squeezing data is done analogously to the recompression data in Section 4.2, splitting into pre and post drop and fitting the three sines to get the parameters of motion in all three axes. We do however apply active post-selection here: runs whose radial amplitudes of fitted frequencies deviate from the first-dataset medians by more than a tolerance (set to 3 median absolute deviations for amplitudes and $0.9 \times 8, \mathrm{kHz}$ for frequencies) are excluded. This is in detail covered in Section A.2.7, and justified as we only care for the squeezing in the axial mode, and both large radial amplitudes and frequency deviations are indicative of runs where the particle has explored nonlinearities (see Equation (2.21) and Figure 4.8) which would lead to a mixing of the modes and a breakdown of the simple independent-mode model.
The data follow the theoretical curve very well, although there seems to be a slight systematic offset as well as experimental difficulties in reaching the squeezing limit. Also with more data points and stronger post-selection, we could not increase the quality of the data, which is why it is not shown here. The minimum near $t_\mathrm{p} \approx 0.36\,\mathrm{\mu s}$ marks the optimal pulse for this choice of τ_1, τ_2. As the $1\,\sigma$ error is still well above both squeezing limits, the data do not allow any reasonable claim on achieving squeezing. Nevertheless, the data show clearly that the model is not wrong and that the procedure could potentially work. For future work, also the timings could be optimized, as well as one could explore squeezing with longer pulses t_p, weaker pulses, full tomography of the state after the protocol, utilizing rotations in phase space to speed up the squeezing or a spin-echo type protocol to mitigate the effect of residual static forces or a mixture of these.

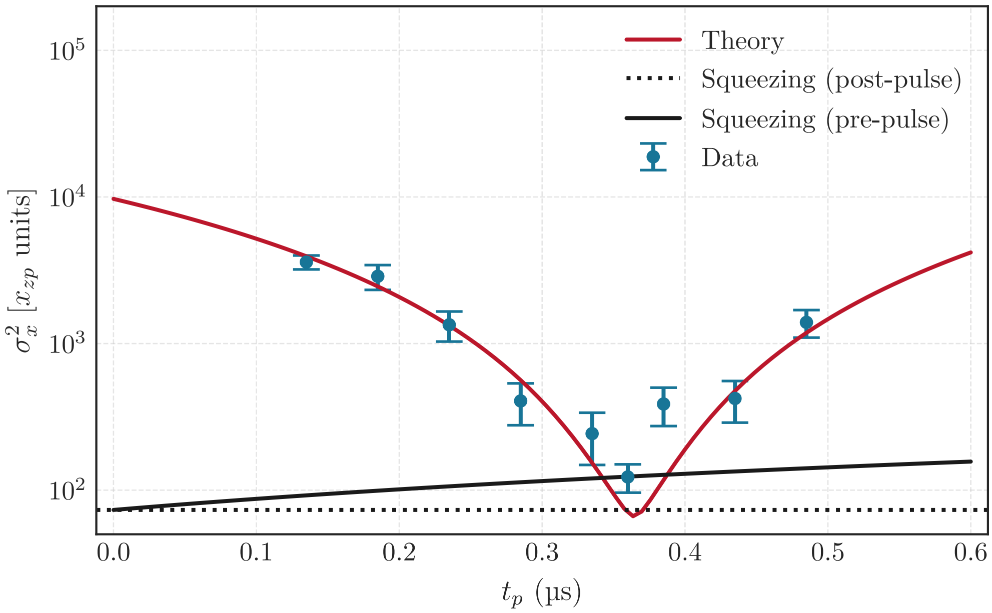

Figure 4.26.: **Squeezing as a function of realignment pulse time.** The maximum ensemble variance max$[\sigma_z^2]$ in units of x_{zp}^2 versus realignment pulse duration t_{p} for $\tau_1 = 20\,\mu\mathrm{s}$, $\tau_2 = 15\,\mu\mathrm{s}$ is shown here. Each point is 150 repetitions, with error bars showing the $1\,\sigma$ confidence interval from bootstrap resampling (see Section A.2.8), but also post selected to exclude runs with large radial amplitudes or frequency deviations). The actual number of runs included in each point is shown in Figure A.5. Red curve: simple Lyapunov-equation theory for a thermal state with $\bar{n} = 12$ and pressure $P \approx 3.9 \times 10^{-8}$ mbar. Black lines: ground-state ($\bar{n} = 0$) squeezing limits, taking into account the rotation pulse (solid) and without (dotted). The minimum near $t_{\mathrm{p}} \approx 0.36\,\mu\mathrm{s}$ marks the optimal pulse for this choice of τ_1, τ_2.

5. Conclusion and Outlooks

With the improvements and methods in this thesis, the static force compensated local free evolution of a charged nanoparticle has reached the 1 ms timescale, moving the technique to a regime only small improvements away from that found in the protocol of [18].
We also demonstrated that the charge state of the environment can be characterised by our compensation scans, is stable, and can be manipulated deterministically via thermionic emission from a filament. While we also showed that after manipulation the resulting environmental charge state then decays, by fixing the required accuracy of the compensation voltages one can predict the timescale after wich the residual drift is negligible, so it still is a useful tool, as the charge state does not return to the initial state.
The squeezing measurement we performed as an addition is not quite enough to claim a successful demonstration of squeezing of the COM motion, but is nevertheless an impressive result given the non-ground-state cooled state we started with.

The experimental setup was mostly there already at the start of the thesis, and all upgrades mentioned were added during the duration of my work on this project. I was involved in characterising the additions — measurements of amplifier nonlinearities, voltage drops, troubleshooting of switching issues, and more. For the compensation matrix I contributed to the conceptual discussions and was responsible for the proper software implementation. I extended the existing control software to also support the various measurement scans, and the data acquisition functionality used throughout this work, along with general bug fixes. All data analysis (unless noted otherwise) presented in this thesis was written in Python from scratch, but adapted from earlier MATLAB-based workflows used for the plots in [1] and for real-time data analysis.

The current limitation on free evolution time is determining the compensation voltages accurately enough, perhaps requiring more iterations of the method done here. Investigating this might require some further improvements on the stability of the DC-level. Potentially also a further measurement like Figure 4.18, zoomed in to properly resolve the voltage minimum as best as possible, could provide insight in fluctuations of the required voltages. High frequency changes to these voltages are harder to detect with this method, as the full compensation scan takes time, limiting the resolution bandwith to slower than minute drifts.
To look into the future, the methods developed are already being used in another setup, where the compensation protocol has been adapted to a different optical trapping setup in a standing wave and has allowed them to reach 200 µs free evolution times

without any major changes needed. As that setup has higher trap frequencies and boasts improved performance due to the optical configuration, these shorter timescales are already interesting and could lead to more squeezing than shown in this thesis.
The setup used here was envisioned to be a testbed for methods and techniques, and the future path likely will be to move on from the not-quite-squeezing efforts to move to improving the protocol sequence. Concepts that are being considered are spin-echo-like sequences to further suppress the effect of residual forces, accelerated expansion schemes aided by rotating the state in phase space between free evolutions, state tomography and more.
To this end, while writing the thesis, I have in parallel adapted the control software to easily support arbitrary sequences of pulses. Using the AOM of the setup, it is also possible to implement weaker pulses, so exploring the interplay of this in all the mentioned ideas is a further interesting path forward. In all of them, the developed compensation protocol will be a key component.

References

[1] D. Steiner *et al.*, 'Free expansion of a charged nanoparticle via electrostatic compensation', *Appl. Phys. Lett.*, vol. 127, no. 19, p. 191 103, 10th Nov. 2025, ISSN: 0003-6951. DOI: `10.1063/5.0303460`. Accessed: 29th Nov. 2025. [Online]. Available: `https://doi.org/10.1063/5.0303460`.

[2] C. Gonzalez-Ballestero, M. Aspelmeyer, L. Novotny, R. Quidant and O. Romero-Isart, 'Levitodynamics: Levitation and control of microscopic objects in vacuum', *Science*, vol. 374, no. 6564, eabg3027, 8th Oct. 2021, ISSN: 0036-8075, 1095-9203. DOI: `10.1126/science.abg3027`. Accessed: 22nd Jan. 2025. [Online]. Available: `https://www.science.org/doi/10.1126/science.abg3027`.

[3] J. Millen, T. S. Monteiro, R. Pettit and A. N. Vamivakas, 'Optomechanics with levitated particles', *Rep. Prog. Phys.*, vol. 83, no. 2, p. 026 401, 1st Feb. 2020, ISSN: 0034-4885, 1361-6633. DOI: `10.1088/1361-6633/ab6100`. Accessed: 12th Sep. 2025. [Online]. Available: `https://iopscience.iop.org/article/10.1088/1361-6633/ab6100`.

[4] D. Rugar, R. Budakian, H. J. Mamin and B. W. Chui, 'Single spin detection by magnetic resonance force microscopy', *Nature*, vol. 430, no. 6997, pp. 329–332, Jul. 2004, ISSN: 1476-4687. DOI: `10.1038/nature02658`. Accessed: 2nd Feb. 2026. [Online]. Available: `https://www.nature.com/articles/nature02658`.

[5] D. G. de Oteyza *et al.*, 'Direct Imaging of Covalent Bond Structure in Single-Molecule Chemical Reactions', *Science*, vol. 340, no. 6139, pp. 1434–1437, 21st Jun. 2013. DOI: `10.1126/science.1238187`. Accessed: 2nd Feb. 2026. [Online]. Available: `https://www.science.org/doi/10.1126/science.1238187`.

[6] LIGO Scientific Collaboration and Virgo Collaboration *et al.*, 'Observation of Gravitational Waves from a Binary Black Hole Merger', *Phys. Rev. Lett.*, vol. 116, no. 6, p. 061 102, 11th Feb. 2016. DOI: `10.1103/PhysRevLett.116.061102`. Accessed: 8th Feb. 2026. [Online]. Available: `https://link.aps.org/doi/10.1103/PhysRevLett.116.061102`.

[7] G. Ranjit, D. P. Atherton, J. H. Stutz, M. Cunningham and A. A. Geraci, 'Attonewton force detection using microspheres in a dual-beam optical trap in high vacuum', *Phys. Rev. A*, vol. 91, no. 5, p. 051 805, 26th May 2015, ISSN: 1050-2947, 1094-1622. DOI: `10.1103/PhysRevA.91.051805`. Accessed: 2nd Feb. 2026. [Online]. Available: `https://link.aps.org/doi/10.1103/PhysRevA.91.051805`.

[8] M. Rossi *et al.*, 'Quantum Delocalization of a Levitated Nanoparticle', *Physical Review Letters*, vol. 135, no. 8, pp. 083 601–083 601, 24th Jun. 2025, ISSN: 0031-9007. DOI: `10.1103/2yzc-fsm3`.

[9] U. Delić *et al.*, 'Cooling of a levitated nanoparticle to the motional quantum ground state', *Science*, vol. 367, no. 6480, pp. 892–895, 21st Feb. 2020. DOI: 10.1126/science.aba3993. Accessed: 29th Jan. 2025. [Online]. Available: https://www.science.org/doi/full/10.1126/science.aba3993.

[10] M. Kamba and K. Aikawa, 'Revealing the Velocity Uncertainties of a Levitated Particle in the Quantum Ground State', *Phys. Rev. Lett.*, vol. 131, no. 18, p. 183 602, 30th Oct. 2023, ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.131.183602. Accessed: 12th Sep. 2025. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.131.183602.

[11] J. Piotrowski *et al.*, 'Simultaneous ground-state cooling of two mechanical modes of a levitated nanoparticle', *Nat. Phys.*, vol. 19, no. 7, pp. 1009–1013, Jul. 2023, ISSN: 1745-2473, 1745-2481. DOI: 10.1038/s41567-023-01956-1. Accessed: 12th Sep. 2025. [Online]. Available: https://www.nature.com/articles/s41567-023-01956-1.

[12] L. Magrini *et al.*, 'Real-time optimal quantum control of mechanical motion at room temperature', *Nature*, vol. 595, no. 7867, pp. 373–377, 15th Jul. 2021, ISSN: 0028-0836, 1476-4687. DOI: 10.1038/s41586-021-03602-3. Accessed: 27th Mar. 2025. [Online]. Available: https://www.nature.com/articles/s41586-021-03602-3.

[13] V. Giovannetti, S. Lloyd and L. Maccone, 'Quantum-Enhanced Measurements: Beating the Standard Quantum Limit', *Science*, vol. 306, no. 5700, pp. 1330–1336, 19th Nov. 2004, ISSN: 0036-8075, 1095-9203. DOI: 10.1126/science.1104149. Accessed: 8th Feb. 2026. [Online]. Available: https://www.science.org/doi/10.1126/science.1104149.

[14] M. Kamba, S. Otabe, K. Funo, T. Sagawa and K. Aikawa. 'A levitated nano-accelerometer sensitized by quantum quench'. arXiv: 2601.21119 [quant-ph], Accessed: 1st Feb. 2026. [Online]. Available: http://arxiv.org/abs/2601.21119, pre-published.

[15] M. Skrabulis *et al.* 'Nanomechanical sensor resolving impulsive forces below its zero-point fluctuations'. arXiv: 2601.19392 [quant-ph], Accessed: 31st Jan. 2026. [Online]. Available: http://arxiv.org/abs/2601.19392, pre-published.

[16] F. Tebbenjohanns, M. L. Mattana, M. Rossi, M. Frimmer and L. Novotny, 'Quantum control of a nanoparticle optically levitated in cryogenic free space', *Nature*, vol. 595, no. 7867, pp. 378–382, 15th Jul. 2021, ISSN: 0028-0836, 1476-4687. DOI: 10.1038/s41586-021-03617-w. Accessed: 27th Mar. 2025. [Online]. Available: https://www.nature.com/articles/s41586-021-03617-w.

[17] S. Pedalino, B. E. Ramírez-Galindo, R. Ferstl, K. Hornberger, M. Arndt and S. Gerlich, 'Probing quantum mechanics with nanoparticle matter-wave interferometry', *Nature*, vol. 649, no. 8098, pp. 866–870, Jan. 2026, ISSN: 1476-4687. DOI: 10.1038/s41586-025-09917-9. Accessed: 21st Jan. 2026. [Online]. Available: https://www.nature.com/articles/s41586-025-09917-9.

[18] L. Neumeier, M. A. Ciampini, O. Romero-Isart, M. Aspelmeyer and N. Kiesel, 'Fast quantum interference of a nanoparticle via optical potential control', *Proc. Natl. Acad. Sci. U.S.A.*, vol. 121, no. 4, e2306953121, 23rd Jan. 2024, ISSN: 0027-8424, 1091-6490. DOI: 10.1073/pnas.2306953121. arXiv: 2207.12539 [cond-mat, physics:quant-ph]. Accessed: 22nd Jan. 2025. [Online]. Available: https://pnas.org/doi/10.1073/pnas.2306953121.

[19] T. Weiss, M. Roda-Llordes, E. Torrontegui, M. Aspelmeyer and O. Romero-Isart, 'Large Quantum Delocalization of a Levitated Nanoparticle Using Optimal Control: Applications for Force Sensing and Entangling via Weak Forces', *Phys. Rev. Lett.*, vol. 127, no. 2, p. 023 601, 7th Jul. 2021, ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.127.023601. Accessed: 24th Jun. 2025. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.127.023601.

[20] A. Bassi, K. Lochan, S. Satin, T. P. Singh and H. Ulbricht, 'Models of wave-function collapse, underlying theories, and experimental tests', *Rev. Mod. Phys.*, vol. 85, no. 2, pp. 471–527, 2nd Apr. 2013. DOI: 10.1103/RevModPhys.85.471. Accessed: 9th Feb. 2026. [Online]. Available: https://link.aps.org/doi/10.1103/RevModPhys.85.471.

[21] D. Carney *et al.*, 'Mechanical quantum sensing in the search for dark matter', *Quantum Sci. Technol.*, vol. 6, no. 2, p. 024 002, 1st Apr. 2021, ISSN: 2058-9565. DOI: 10.1088/2058-9565/abcfcd. Accessed: 18th Oct. 2025. [Online]. Available: https://iopscience.iop.org/article/10.1088/2058-9565/abcfcd.

[22] I. S. Chowdhury, B. P. Akhouri, S. Haque, M. H. Bacci and E. Howard, 'Search for Dark Matter Using Levitated Nanoparticles Within a Bessel–Gaussian Beam via Yukawa Coupling', *Sensors (Basel)*, vol. 25, no. 23, p. 7138, 22nd Nov. 2025, ISSN: 1424-8220. DOI: 10.3390/s25237138. PMID: 41374513. Accessed: 5th Feb. 2026. [Online]. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC12694501/.

[23] O. Romero-Isart, 'Coherent inflation for large quantum superpositions of levitated microspheres', *New J. Phys.*, vol. 19, no. 12, p. 123 029, 13th Dec. 2017, ISSN: 1367-2630. DOI: 10.1088/1367-2630/aa99bf. Accessed: 12th Sep. 2025. [Online]. Available: https://iopscience.iop.org/article/10.1088/1367-2630/aa99bf.

[24] G. F. M. Tomassi *et al.*, 'Accelerated state expansion of a nanoparticle in a dark inverted potential', *Phys. Rev. Research*, vol. 8, no. 1, p. L012026, 26th Jan. 2026, ISSN: 2643-1564. DOI: 10.1103/r7t1-w311. Accessed: 1st Apr. 2026. [Online]. Available: https://link.aps.org/doi/10.1103/r7t1-w311.

[25] S. Dago *et al.*, 'Stabilizing nanoparticles in the intensity minimum: Feedback levitation on an inverted potential', *Opt. Express*, vol. 32, no. 25, p. 45 133, 2nd Dec. 2024, ISSN: 1094-4087. DOI: 10.1364/OE.541267. Accessed: 12th Sep. 2025. [Online]. Available: https://opg.optica.org/abstract.cfm?URI=oe-32-25-45133.

[26] V. Mlynář *et al.*, 'Feedback stabilization of a nanoparticle at the intensity minimum of an optical double-well potential', *Control Engineering Practice*, vol. 168, p. 106 665, 1st Mar. 2026, ISSN: 0967-0661. DOI: 10.1016/j.conengprac.2025.106665. Accessed: 27th Dec. 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0967066125004277.

[27] M. Kamba, N. Hara and K. Aikawa, 'Quantum squeezing of a levitated nanomechanical oscillator', *Science*, vol. 389, no. 6766, pp. 1225–1228, 18th Sep. 2025. DOI: 10.1126/science.ady4652. Accessed: 1st Apr. 2026. [Online]. Available: https://www.science.org/doi/abs/10.1126/science.ady4652.

[28] E. Bonvin *et al.*, 'State Expansion of a Levitated Nanoparticle in a Dark Harmonic Potential', *Phys. Rev. Lett.*, vol. 132, no. 25, p. 253 602, 21st Jun. 2024, ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.132.253602. Accessed: 12th Sep. 2025. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.132.253602.

[29] S. Otabe, M. Kamba, Y. Kojima and K. Aikawa. 'Time-of-flight force sensing below the quantum zero-point fluctuation'. arXiv: 2605.09854 [quant-ph], Accessed: 12th May 2026. [Online]. Available: http://arxiv.org/abs/2605.09854, prepublished.

[30] E. Hebestreit, M. Frimmer, R. Reimann and L. Novotny, 'Sensing Static Forces with Free-Falling Nanoparticles', *Phys. Rev. Lett.*, vol. 121, no. 6, p. 063 602, 7th Aug. 2018, ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.121.063602.

[31] M. L. Mattana *et al.*, 'Trap-to-trap free falls with an optically levitated nanoparticle', *Phys. Rev. A*, vol. 113, no. 2, p. 023 510, 9th Feb. 2026. DOI: 10.1103/47pz-74fc. Accessed: 20th May 2026. [Online]. Available: https://link.aps.org/doi/10.1103/47pz-74fc.

[32] G. Prakash, S. Herrmann, R. B. Bergmann and C. Vogt, 'Release and recapture of silica nanoparticles from an optical trap in weightlessness', *npj Microgravity*, vol. 12, no. 1, p. 37, 27th Apr. 2026, ISSN: 2373-8065. DOI: 10.1038/s41526-026-00596-y. Accessed: 20th May 2026. [Online]. Available: https://www.nature.com/articles/s41526-026-00596-y.

[33] G. Gasbarri *et al.*, 'Testing the foundation of quantum physics in space via Interferometric and non-interferometric experiments with mesoscopic nanoparticles', *Commun Phys*, vol. 4, no. 1, p. 155, 7th Jul. 2021, ISSN: 2399-3650. DOI: 10.1038/s42005-021-00656-7. Accessed: 12th Sep. 2025. [Online]. Available: https://www.nature.com/articles/s42005-021-00656-7.

[34] A. Ashkin, 'Acceleration and Trapping of Particles by Radiation Pressure', *Phys. Rev. Lett.*, vol. 24, no. 4, pp. 156–159, 26th Jan. 1970, ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.24.156. Accessed: 15th Jun. 2025. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.24.156.

[35] A. Ashkin, J. M. Dziedzic, J. E. Bjorkholm and S. Chu, 'Observation of a single-beam gradient force optical trap for dielectric particles',

[36] H. C. Hulst and H. C. van de Hulst, *Light Scattering by Small Particles*. Courier Corporation, 1st Jan. 1981, 502 pp., ISBN: 978-0-486-64228-4. Google Books: PlHfPMVAFRcC.

[37] P. Lebedew, 'Untersuchungen über die Druckkräfte des Lichtes', *Annalen der Physik*, vol. 311, no. 11, pp. 433–458, Jan. 1901, ISSN: 0003-3804, 1521-3889. DOI: 10.1002/andp.19013111102. Accessed: 27th Sep. 2024. [Online]. Available: https://onlinelibrary.wiley.com/doi/10.1002/andp.19013111102.

[38] E. F. Nichols and G. F. Hull, 'The Pressure Due to Radiation', *Proceedings of the American Academy of Arts and Sciences*, vol. 38, no. 20, p. 559, 1903, ISSN: 0199-9818. DOI: 10.2307/20021808. JSTOR: 10.2307/20021808. Accessed: 17th Sep. 2024. [Online]. Available: https://www.jstor.org/stable/10.2307/20021808?origin=crossref.

[39] L. Novotny and B. Hecht, *Principles of Nano-Optics*, 3rd ed. Cambridge: Cambridge University Press, 2025, ISBN: 978-1-108-47894-6. DOI: 10.1017/9781108781503. Accessed: 15th Feb. 2026. [Online]. Available: https://www.cambridge.org/core/books/principles-of-nanooptics/D476A9D8F068D42EEB51FEF944CD76CB.

[40] L. Novotny and B. Hecht, *Principles of Nano-Optics*, 3rd ed. Cambridge University Press, 6th Sep. 2012, 583 pp., ISBN: 978-1-139-56045-0. Google Books: RHC_AwAAQBAJ.

[41] V. Jain, J. Gieseler, C. Moritz, C. Dellago, R. Quidant and L. Novotny, 'Direct Measurement of Photon Recoil from a Levitated Nanoparticle', *Phys. Rev. Lett.*, vol. 116, no. 24, p. 243 601, 13th Jun. 2016. DOI: 10.1103/PhysRevLett.116.243601. Accessed: 1st Feb. 2026. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.116.243601.

[42] S. Earnshaw, 'On the Nature of the Molecular Forces which Regulate the Constitution of the Luminiferous Ether', *Transactions of the Cambridge Philosophical Society*, vol. 7, p. 97, 1st Jan. 1848. Accessed: 28th Mar. 2026. [Online]. Available: https://ui.adsabs.harvard.edu/abs/1848TCaPS...7...97E.

[43] T. W. Anderson, *An Introduction to Multivariate Statistical Analysis* (Wiley Series in Probability and Statistics), 3rd ed. Hoboken, N.J: Wiley-Interscience, 2003, 721 pp., ISBN: 978-0-471-36091-9.

[44] A. Serafini, *Quantum Continuous Variables: A Primer of Theoretical Methods*, 2nd ed. Boca Raton: CRC Press, 1st Jun. 2023, ISBN: 978-1-003-25097-5. DOI: 10.1201/9781003250975. Accessed: 22nd Mar. 2026. [Online]. Available: https://www.taylorfrancis.com/books/9781003250975.

[45] M. G. Genoni, L. Lami and A. Serafini, 'Conditional and unconditional Gaussian quantum dynamics', *Contemporary Physics*, vol. 57, no. 3, pp. 331–349, 2nd Jul. 2016, ISSN: 0010-7514, 1366-5812. DOI: 10.1080/00107514.2015.1125624. Accessed: 12th Sep. 2025. [Online]. Available: http://www.tandfonline.com/doi/full/10.1080/00107514.2015.1125624.

[46] F. Ricci, R. Quidant and R. Rica Alarcón, 'Levitodynamics toward force nanosensors in vacuum', Universitat Politècnica de Catalunya, 22nd Feb. 2019. DOI: `10.5821/dissertation-2117-131414`. HDL: `2117/131414`. Accessed: 11th Jun. 2024. [Online]. Available: `http://hdl.handle.net/2117/131414`.

[47] H. Uys, J. D. Perreault and A. D. Cronin, 'Matter-Wave Decoherence due to a Gas Environment in an Atom Interferometer', *Phys. Rev. Lett.*, vol. 95, no. 15, p. 150 403, 5th Oct. 2005. DOI: `10.1103/PhysRevLett.95.150403`. Accessed: 9th May 2026. [Online]. Available: `https://link.aps.org/doi/10.1103/PhysRevLett.95.150403`.

[48] F. van der Laan, 'Rotational Levitodynamics', 2022. DOI: `10.3929/ETHZ-B-000591484`. HDL: `20.500.11850/591484`. Accessed: 21st Jan. 2026. [Online]. Available: `http://hdl.handle.net/20.500.11850/591484`.

[49] C. Bach, 'Detection of a levitated particle's mechanical motion close to the Heisenberg limit', Universität Wien, Vienna, 2021.

[50] F. Tebbenjohanns, M. Frimmer and L. Novotny, 'Optimal position detection of a dipolar scatterer in a focused field', *Phys. Rev. A*, vol. 100, no. 4, p. 043 821, 14th Oct. 2019. DOI: `10.1103/PhysRevA.100.043821`. Accessed: 30th Nov. 2025. [Online]. Available: `https://link.aps.org/doi/10.1103/PhysRevA.100.043821`.

[51] *Phase-locked loop*, in *Wikipedia*, 20th Mar. 2026. Accessed: 16th Apr. 2026. [Online]. Available: `https://en.wikipedia.org/w/index.php?title=Phase-locked_loop&oldid=1344455864#Clock_analogy`.

[52] N. Kuznetsov, G. Leonov, M. Yuldashev and R. Yuldashev, 'Rigorous mathematical definitions of the hold-in and pull-in ranges for phase-locked loops', *IFAC-PapersOnLine*, vol. 48, no. 11, pp. 710–713, 2015, ISSN: 24058963. DOI: `10.1016/j.ifacol.2015.09.272`. Accessed: 16th Apr. 2026. [Online]. Available: `https://linkinghub.elsevier.com/retrieve/pii/S2405896315013531`.

[53] G. A. Leonov, N. V. Kuznetsov, M. V. Yuldashev and R. V. Yuldashev, 'Hold-In, Pull-In, and Lock-In Ranges of PLL Circuits: Rigorous Mathematical Definitions and Limitations of Classical Theory', *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 62, no. 10, pp. 2454–2464, Oct. 2015, ISSN: 1558-0806. DOI: `10.1109/TCSI.2015.2476295`. Accessed: 16th Apr. 2026. [Online]. Available: `https://ieeexplore.ieee.org/document/7277189`.

[54] C. P. Slichter, *Principles of Magnetic Resonance* (Springer Series in Solid-State Sciences 1), 3., enlarged and updated ed. Berlin [u.a.]: Springer, 1990, ISBN: 978-3-540-50157-2.

[55] M. Frimmer, K. Luszcz, S. Ferreiro, V. Jain, E. Hebestreit and L. Novotny, 'Controlling the net charge on a nanoparticle optically levitated in vacuum', *Phys. Rev. A*, vol. 95, no. 6, p. 061 801, 6th Jun. 2017, ISSN: 2469-9926, 2469-9934. DOI: `10.1103/PhysRevA.95.061801`. Accessed: 19th Mar. 2025. [Online]. Available: `http://link.aps.org/doi/10.1103/PhysRevA.95.061801`.

[56] D. C. Moore, A. D. Rider and G. Gratta, 'Search for Millicharged Particles Using Optically Levitated Microspheres', *Phys. Rev. Lett.*, vol. 113, no. 25, p. 251 801, 16th Dec. 2014, ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.113.251801. Accessed: 12th Sep. 2025. [Online]. Available: https://link.aps.org/doi/10.1103/PhysRevLett.113.251801.

[57] H. Inaba, T. Ohmi, T. Yoshida and T. Okada, 'Neutralization of static electricity by soft X-rays and vacuum UV radiation', *Journal of Electrostatics*, vol. 33, no. 1, pp. 15–42, Jun. 1994, ISSN: 03043886. DOI: 10.1016/0304-3886(94)90061-2. Accessed: 19th Mar. 2025. [Online]. Available: https://linkinghub.elsevier.com/retrieve/pii/0304388694900612.

[58] L. Zhu, Y. Yang, J. Cai, X. Xu, L. Ma and J. Luo. 'Effect of UV light irradiation on charge neutralization in XPS measurements'. version 2, Accessed: 19th Mar. 2025. [Online]. Available: https://arxiv.org/abs/2409.00663, pre-published.

[59] T. Westphal *et al.*, 'Measurement of gravitational coupling between millimetre-sized masses.', *Nature*, vol. 591, no. 7849, pp. 225–228, 11th Mar. 2021, ISSN: 0028-0836, 1476-4687. DOI: 10.1038/s41586-021-03250-7. PMID: 33692556. [Online]. Available: https://www.nature.com/articles/s41586-021-03250-7.

[60] R. Weiss, 'Charging of the test masses past, present and future',

[61] D. Ugolini, Q. Funk and T. Amen, 'Note: Discharging fused silica test masses with ionized nitrogen', *Review of Scientific Instruments*, vol. 82, no. 4, p. 046 108, 1st Apr. 2011, ISSN: 0034-6748, 1089-7623. DOI: 10.1063/1.3579500. Accessed: 12th Sep. 2025. [Online]. Available: https://pubs.aip.org/rsi/article/82/4/046108/353543/Note-Discharging-fused-silica-test-masses-with.

[62] E. Hebestreit, M. Frimmer, R. Reimann, C. Dellago, F. Ricci and L. Novotny, 'Calibration and temperature measurement of levitated optomechanical sensors', *Review of Scientific Instruments*, vol. 89, no. 3, p. 033 111, 1st Mar. 2018, ISSN: 0034-6748, 1089-7623. DOI: 10.1063/1.5017119. arXiv: 1711.09049 [physics]. Accessed: 27th Mar. 2025. [Online]. Available: http://arxiv.org/abs/1711.09049.

[63] C. R. Harris *et al.*, 'Array programming with NumPy', *Nature*, vol. 585, no. 7825, pp. 357–362, 17th Sep. 2020, ISSN: 0028-0836, 1476-4687. DOI: 10.1038/s41586-020-2649-2. Accessed: 18th May 2026. [Online]. Available: https://www.nature.com/articles/s41586-020-2649-2.

[64] P. Virtanen *et al.*, 'SciPy 1.0: Fundamental algorithms for scientific computing in Python', *Nat Methods*, vol. 17, no. 3, pp. 261–272, 2nd Mar. 2020, ISSN: 1548-7091, 1548-7105. DOI: 10.1038/s41592-019-0686-2. Accessed: 17th Sep. 2024. [Online]. Available: https://www.nature.com/articles/s41592-019-0686-2.

[65] J. D. Hunter, 'Matplotlib: A 2D Graphics Environment', *Computing in Science & Engineering*, vol. 9, no. 3, pp. 90–95, May 2007, ISSN: 1558-366X. DOI: 10.1109/MCSE.2007.55. Accessed: 18th May 2026. [Online]. Available: https://ieeexplore.ieee.org/document/4160265.

[66] M. L. Waskom, 'Seaborn: Statistical data visualization', *Journal of Open Source Software*, vol. 6, no. 60, p. 3021, 6th Apr. 2021, ISSN: 2475-9066. DOI: `10.21105/joss.03021`. Accessed: 18th May 2026. [Online]. Available: `https://joss.theoj.org/papers/10.21105/joss.03021`.

[67] B. Efron, 'Bootstrap Methods: Another Look at the Jackknife', *The Annals of Statistics*, vol. 7, no. 1, pp. 1–26, Jan. 1979, ISSN: 0090-5364, 2168-8966. DOI: `10.1214/aos/1176344552`. Accessed: 28th Apr. 2026. [Online]. Available: `https://projecteuclid.org/journals/annals-of-statistics/volume-7/issue-1/Bootstrap-Methods-Another-Look-at-the-Jackknife/10.1214/aos/1176344552.full`.

[68] A. H. Al-Mohy and N. J. Higham, 'A New Scaling and Squaring Algorithm for the Matrix Exponential', *SIAM J. Matrix Anal. & Appl.*, vol. 31, no. 3, pp. 970–989, Jan. 2010, ISSN: 0895-4798, 1095-7162. DOI: `10.1137/09074721X`. Accessed: 4th May 2026. [Online]. Available: `http://epubs.siam.org/doi/10.1137/09074721X`.

A. Appendix

In this section some of the code was shortened using coding-specific large language models to only include the relevant parts for the analysis, as once printed, lines of code offer no great reusability in any case. The full code is available upon reasonable request as a digital file.

A.1. Detector Calibration

The purpose of detector calibration is to determine the conversion factor between the measured detector signal (in volts) and the actual physical displacement of the nanoparticle (in metres). This is achieved by analysing the thermal motion of the particle at high gas pressure (e.g., 9 mbar) and applying the equipartition theorem. We follow closely [62].

At high pressures, the nanoparticle is in thermal equilibrium with the surrounding gas due to frequent collisions. The particle's motion in the harmonic trapping potential can be described as a damped harmonic oscillator. According to the equipartition theorem, the average potential energy of the particle in one dimension is given by:

$$\langle U \rangle = \frac{1}{2} m \Omega_0^2 \langle x^2 \rangle = \frac{1}{2} k_B T \tag{A.1}$$

where m is the particle's mass, Ω_0 is the angular trapping frequency, $\langle x^2 \rangle$ is the mean squared displacement from the trap centre, k_B is the Boltzmann constant, and T is the gas temperature.

From this, we can express the mean squared displacement as:

$$\langle x^2 \rangle = \frac{k_B T}{m \Omega_0^2} \tag{A.2}$$

The detector measures a voltage signal V_{sig} proportional to the particle's displacement x ($V_{\mathrm{sig}} = \alpha x$), where α is the desired calibration factor (in V/m). The mean squared voltage signal, calculated from the measured time trace, is therefore related to the mean squared displacement by:

$$\langle V_{\mathrm{sig}}^2 \rangle = \alpha^2 \langle x^2 \rangle \tag{A.3}$$

By combining these equations, we can solve for the calibration factor α:

$$\alpha = \sqrt{\frac{\langle V_{\mathrm{sig}}^2 \rangle}{\langle x^2 \rangle}} = \sqrt{\frac{\langle V_{\mathrm{sig}}^2 \rangle m \Omega_0^2}{k_B T}} \tag{A.4}$$

By measuring the mean squared voltage from a time trace at a known temperature and pressure, and with prior knowledge of the particle's mass and trapping frequency, we can determine the calibration factor α.

There are now several methods to actually get a value for this α. We can fit a Lorentzian (which is the shape expected for the displacement power spectral density of a damped harmonic oscillator) to the power spectral density (PSD) of the measured time trace, leaving the pressure p, the frequency Ω_0, and the calibration factor α as free parameters. By then cross-checking the pressure and constraining the frequency, we can be confident that the fitted α is correct. For the fitted data in Figure A.1, we get a calibration factor of $\alpha = 28.23\,\mathrm{mV\,nm^{-1}}$.

Additionally, we can bandpass the timetrace data around the trapping frequency with a 30 kHz Butterworth filter and then calculate the variance of this filtered signal, which, using Equation (A.4), can be directly converted to the calibration factor $\alpha = 25.97\,\mathrm{mV\,nm^{-1}}$.

Parseval's theorem states that the total power should be equal in the time or frequency domain, so we can also integrate the PSD over a frequency range that cuts out the noise and contribution of other modes (using the same area that is also used for the fits) and then use this with Equation (A.4) to get $\alpha = 26.40\,\mathrm{mV\,nm^{-1}}$.

The reported α in Figure A.1 is then the average of these three methods, since there is no clear indication that any should be preferred and they all agree well.

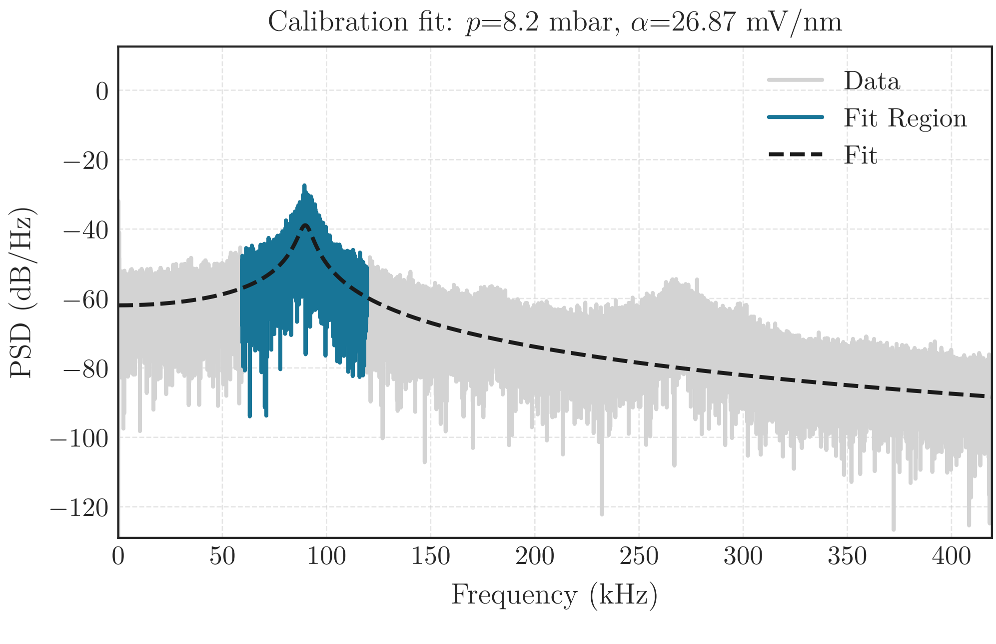

Figure A.1.: **Example plot of the calibration procedure.** Shown is the detector signal PSD (grey), with a Lorentzian fit (dashed black) to the region of interest marked in blue. The area marked is also used for the PSD integration method. The reported pressure comes from the Lorentzian fit, and the calibration factor is the average of the three methods described in the text.

Furthermore, this calibration can be extended to different experimental conditions. The detector signal is proportional to the local oscillator (LO) field amplitude used for detection. Therefore, the calibration factor α is expected to scale with the square root of the LO power as $V_{\text{sig}} \propto \sqrt{P_{\text{LO}}}$.
This relationship allows for extrapolation of the calibration to higher LO power levels (often used in later measurements) by performing the calibration at a low LO power and then using a power meter to measure the relative increase in LO power for other measurements.
This was confirmed with measurements at different LO power levels, showing that the calibration factor α indeed scales with the square root of the LO power, as expected, and is visually represented in Figure A.2.

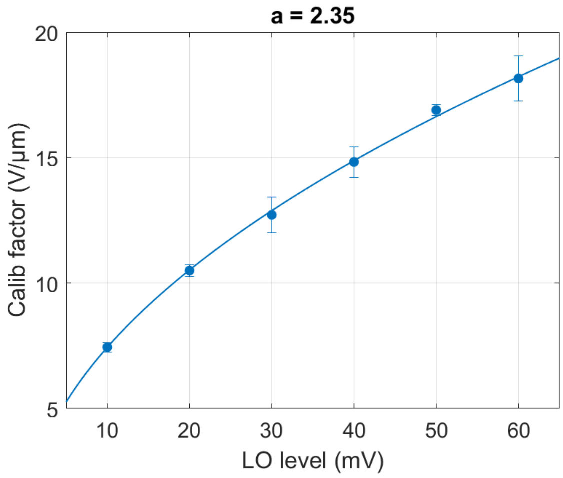

Figure A.2.: **Detector calibration factor scaling with local oscillator power.** The calibration factor α is measured at different LO power levels, showing a clear square root dependence as expected from the theory of homodyne detection. The solid line represents a fit to the data using the function $\alpha(P_{\text{LO}}) = \alpha_0\sqrt{P_{\text{LO}}/P_0}$, where α_0 is the calibration factor at a reference power P_0. Each data point represents 6 calibrations at the LO power measured with a multimeter. Figure credit: Y. Y. Fein

By then measuring the new LO power LO_{new} against the calibration reference power LO_{ref}, the calibration factor can be obtained as $\alpha_{\text{new}} = \alpha_{\text{ref}}\sqrt{LO_{\text{new}}/LO_{\text{ref}}}$. Measuring the LO power and the calibration at the beginning of each measurement is still needed, as, while the scaling holds, the absolute calibration factor can drift for different conditions due to alignment changes, particle size variance, or other factors.

A.2. Data Analysis Methods

The code to produce the plots and analysis in the main text is available upon request, and the following sections and snippets are here to merely give an idea of how the analysis was done and to provide some code snippets for the more intricate parts of the analysis. The code is written in Python and uses standard scientific libraries such as `NumPy`[63], `SciPy`[64], and `Matplotlib`[65] + `Seaborn`[66] for data processing and visualisation.

A.2.1. Moving variance

To calculate a moving variance of the time trace, we use a standard centred moving window approach, making this implementation close to the MATLAB `movvar` function with default settings. `data` is the input time trace and `window_size` is the number of data points in the moving window. The function returns an array of the same length as the input, where each point is the variance of the data within the window centred around that point. The edges are padded with NaNs to maintain the same length as the input, but this means that these NaNs need to be taken care of.

```python
def moving_variance(data, window_size):
    """
    Compute moving variance of a 1D array.

    Parameters
    ----------
    data : array_like
        Input signal.
    window_size : int
        Size of the moving window.

    Returns
    -------
    var : ndarray
        Moving variance. Padded with NaNs at the edges to match input length.
        Matches MATLAB's movvar (default center alignment).
    """
    return pd.Series(data).rolling(window=window_size, center=True).var().to_numpy()
```

Listing A.1: Calculating the moving variance of the time trace.

A.2.2. Processing of a single recompression scan

To parallelise the processing of the recompression scans, we define this worker function, so that multiple bins can be processed in parallel. It then saves the analysis output to an array to be further processed. This process involves the following steps: We first load the file, then map the index to the time, separating the pre- and post-drop data, then we fit the post drop data with a triple sine wave with a linear term `fit_triple_sine_linear` to account for drifts and the radial modes, which gives us a first estimate of the axial frequency. If this fit fails, we can also get a frequency guess from the spectrum of the post drop data. With this frequency guess, we then do another refined fit of the post drop data with a single sine wave with a linear term, which gives us a more stable amplitude estimate. Finally, we fit the pre drop data with a single sine wave with the frequency fixed to the refined post drop fit, to extract the amplitude of the pre drop oscillation, which is related to the energy of the particle before the drop. The code for this process is shown in Listing A.2.

```python
def process_bin_task(args):
    try:
        i, j, fname, directory, drop_val, dt, buffers, t_pre, t_post, freq_guess_rad = args

        dat = np.fromfile(os.path.join(directory, fname, f"{fname}_{j}.bin"), dtype=np.float64)
        N = int(dat.size)

        pre_buffer, post_buffer = buffers
        pre_ind         = int(round(t_pre       / dt))
        post_ind        = int(round(t_post      / dt))
        drop_ind        = int(round(drop_val    / dt))
        pre_buffer_ind  = int(round(pre_buffer  / dt))
        post_buffer_ind = int(round(post_buffer / dt))

        trig      = int(np.floor(N / 2.0 + 0.5))
        ind_start = trig - pre_ind
        ind_end   = trig + post_ind + post_buffer_ind + drop_ind

        sig_plot = dat[ind_start - 1:ind_end]
        x1       = sig_plot[:pre_ind + 1 - pre_buffer_ind]
        x2_start = pre_ind + drop_ind + post_buffer_ind
        x2       = sig_plot[x2_start:]

        t_dat  = np.arange(N) * dt
        t_plot = t_dat[ind_start - 1:ind_end] - t_dat[trig - 1]
        t1 = t_plot[:x1.size]
        t2 = t_plot[x2_start:x2_start + x2.size]

        # Pass 1: 3-mode post-drop fit for axial frequency estimate
        popt2, _ = af.fit_triple_sine_linear(t2, x2, freq_guess=np.asarray(freq_guess_rad))
        if popt2 is None:  # fallback: estimate frequencies from spectrum
            popt2, _ = af.fit_triple_sine_linear(t2, x2,
                freq_guess=_get_frequency_guesses_from_spectrum(x2, dt))

        # Pass 2: refined single-mode post-drop fit
        def post_ref_model(t, A, w, phi, off, slope):
            return A * np.sin(w * t + phi) + slope * t + off

        w0   = freq_guess_rad[0]
        A0   = popt2[0] if popt2 is not None else np.sqrt(2.0 * np.var(x2))
        off0 = popt2[9] if popt2 is not None else np.mean(x2)
        try:
            popt2_ref, pcov2_ref = curve_fit(post_ref_model, t2, x2,
                p0=[A0, w0, 0.0, off0, 0.0],
                bounds=([0, w0 - 2*np.pi*5e3, 0.0, -4.0, -4.0],
```

```
                        [4, w0 + 2*np.pi*5e3, 2*np.pi, 4.0,  4.0]),
                maxfev=100000)
        except RuntimeError:
            popt2_ref, pcov2_ref = None, None

        # Pass 3: pre-drop fit with frequency fixed from Pass 2
        w_pre   = popt2_ref[1] if popt2_ref is not None else freq_guess_rad[0]
        off_pre = float(np.mean(x1))
        pre_model = lambda t, A, phi: A * np.sin(w_pre * t + phi) + off_pre
        try:
            p_pre, _ = curve_fit(pre_model, t1, x1,
                p0=[np.sqrt(2.0 * np.var(x1)), np.pi],
                bounds=([0, 0.0], [4.0, 2*np.pi]), maxfev=50000)
            popt1 = np.array([p_pre[0], p_pre[1], off_pre, 0.0], dtype=float)
        except RuntimeError:
            popt1 = None

        return {
            "indices":   (i, j),
            "popt2":     popt2,     # 3-mode post fit
            "popt2_ref": popt2_ref, # refined post fit [A, w, phi, offset, slope]
            "pcov2_ref": pcov2_ref,
            "popt1":     popt1,     # pre fit [A, phi, offset, slope]
            "trace": {"t1": t1, "x1": x1, "t2": t2, "x2": x2} if j == 0 else None,
        }
    except Exception as err:
        return {"indices": (args[0], args[1]), "popt2": None, "popt2_ref": None,
                "pcov2_ref": None, "popt1": None, "trace": None, "error": str(err)}
```

Listing A.2: Recompression worker function.

A.2.3. Ensemble processing for recompression scans

To check the validity of the fits, we can look at a histogram of the fitted amplitudes for the pre- and post-drop fits, shown in Figure A.3.

From all the fitted parameters, we can then reconstruct sinusoidal trajectories for each drop and combine them to get the ensemble evolution of the variance, leading to the plot shown in Figure A.6 and Figure 4.5. For imperfect recompression, we get a time-varying ensemble variance.

A.2.4. Stable fitting of V_{opt}

Used for Figures 4.8, 4.12 and 4.18 to extract the optimal compensation voltage from parabolic fits.

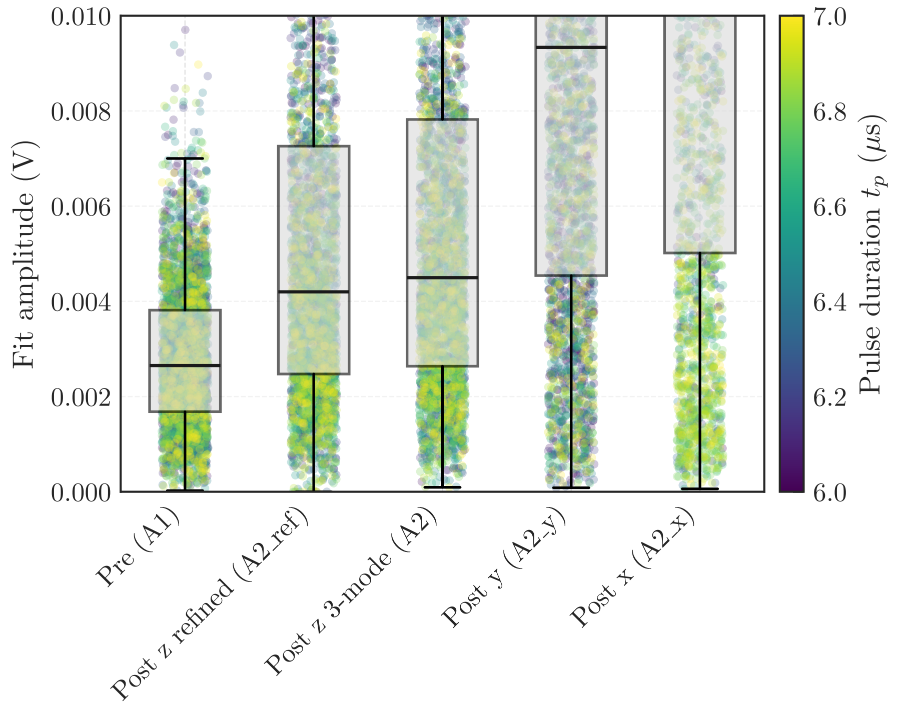

Figure A.3.: **Histogram of fitted amplitudes for pre- and post-drop fits of recompression scans.** The data shown here is the fitted amplitude from the indicated fitting step in Listing A.2. The colour indicates the corresponding pulse duration bin the timetrace belongs in. The data is cut off at 0.01 V, as these high amplitudes could be physical and the fit usually fails by fitting a very small amplitude as this also minimises the loss function used to a local minimum. By seeing a spread in amplitudes and no clustering at low amplitudes for all pulse times, we can conclude that most fits worked well.

The basic method is fitting the data to a parabola, extracting the vertex and plotting it over time. The parabola is fitted to the data using scipy's `least_squares` function, which uses non-linear least squares to fit a function to the data. While the standard `scipy.optimize` might suffice if one takes very good care of sketching out both flanks of the parabola, the `least_squares` function can be made highly robust against outliers and can take a variety of loss functions, which makes it ideal for this task, as it is not a priori clear where the V_{opt} will be when taking the data, so some datasets (especially for the stability scan in Section 4.7.2) might only have one or two datapoints on one side of the flank for fitting the parabola.

```
def fit_parabola_variance(V, relVar, relVar_err, ...):
    parabola = lambda x, a, b, d: a * (x - b)**2 + d
    weighted_residuals = lambda p, x, y, s: (parabola(x, *p) - y) / s

    # Remove invalid points; constrain vertex b to the scan range
    mask = np.isfinite(relVar) & np.isfinite(relVar_err) & (relVar_err > 0)
    bounds = ([0, np.min(V[mask]), -np.inf], [np.inf, np.max(V[mask]), np.inf])
    p0 = [scale_fit_guess, cent_fit_guess, np.min(relVar[mask])]

    # Robust fit with soft-L1 loss to down-weight outliers
    res = least_squares(weighted_residuals, x0=p0,
                        args=(V[mask], relVar[mask], relVar_err[mask]),
                        loss='soft_l1', bounds=bounds)
    pcov = np.linalg.inv(res.jac.T @ res.jac)  # covariance from Jacobian
    ...
```

Listing A.3: Fitting a parabola to the variance data to extract the optimal voltage.

A.2.5. Extraction of mean energy from timetrace

For each drop time τ, the script processes $N = 150$ experimental realisations. Each realisation is stored as a 64-bit floating-point binary file recorded at $f_s = 25\,\text{MHz}$, containing approximately 10 ms of the homodyne position signal with the trigger located at the centre sample. The file is loaded and sliced into a pre-drop segment $x_1(t)$, ending just before the trigger, and a post-drop segment $x_2(t)$, beginning 3 µs after recapture to allow transients from the free evolution to decay.

Triple-sine fitting for the freefall τ-scan

The nanoparticle's homodyne position signal contains oscillations at all three mechanical trap frequencies simultaneously: the axial mode near 91 kHz and the two radial modes near 271 kHz and 305 kHz. Fitting only the axial mode while ignoring the radial contributions would leave correlated residuals and systematically bias the extracted amplitude. The post-drop signal is therefore fitted to a sum of three sine waves with a common linear drift term,

$$x_2(t) = \sum_{k=1}^{3} A_k \sin(\omega_k t + \varphi_k) + c_0 + c_1 t, \tag{A.5}$$

where $\{A_k, \omega_k, \varphi_k\}$ are the amplitude, angular frequency, and phase of each mode, and the linear term $c_0 + c_1 t$ absorbs slow DC drift that is not attributable to mechanical motion.

The 11-parameter nonlinear least-squares problem in Equation (A.5) is sensitive to the quality of the initial frequency guesses: a poor starting point easily causes the solver to settle in a local minimum, usually being an unrealistically low amplitude for the oscillations. To obtain reliable guesses for each individual shot, a longer 4 ms window of the same file is loaded first and its power spectral density is computed using a boxcar periodogram. Spectral peaks are identified in two frequency bands: the axial band (85–95) kHz and the combined radial band (260–300) kHz and the dominant peak in each band is used as the initial guess for ω_k. If no peak is found in a given band, a fixed default is substituted. But from trial and error, we found that the PSD is generally better than always using the default peak. The code for this step is shown in Listing A.4.

```python
def _get_frequency_guesses_from_spectrum(x, dt, num_freqs=3):
    fs = 1.0 / dt
    f, Pxx = af.calculate_psd(x, fs, window='boxcar')
    log_Pxx = np.log10(Pxx + 1e-12)
    peaks, _ = signal.find_peaks(log_Pxx, prominence=0.5)

    band_axial  = (f[peaks] > 85e3)  & (f[peaks] < 95e3)
    band_radial = (f[peaks] > 260e3) & (f[peaks] < 300e3)

    freqs = []
    # Axial: dominant peak in 85-95 kHz
    axial_peaks = peaks[band_axial]
    if len(axial_peaks) > 0:
        freqs.append(f[axial_peaks[np.argmax(Pxx[axial_peaks])]])
    # Radial: top-2 peaks in 260-300 kHz
    radial_peaks = peaks[band_radial]
    if len(radial_peaks) > 0:
        top2 = np.argsort(-Pxx[radial_peaks])[:min(2, len(radial_peaks))]
        freqs.extend(np.sort(f[radial_peaks[top2]]))
    # Fill missing modes with defaults
    defaults = np.array([90e3, 271e3, 305e3])
    while len(freqs) < num_freqs:
        freqs.append(defaults[len(freqs)])
    return np.array(freqs[:num_freqs]) * 2 * np.pi   # rad/s
```

Listing A.4: Spectral frequency initialisation for the triple-sine fit. Run once per file on a long window before the main fit.

Before passing data to `scipy.optimize.curve_fit`, both the time axis and the signal amplitude are normalised. This rescaling brings all fit parameters to $\mathcal{O}(1)$, which improves the fit. After the fit, all parameters are unscaled back to physical units before being stored.

A single call to the solver with wide bounds is not sufficient to reliably resolve all three modes. A narrow band of bounds is equally unreliable, as it is highly dependent on the initial guess and the boundary conditions. We can work around this issue by iteratively adjusting the bounds, each run improving the parameter bounds around the solution of the previous one:

1. **Wide-bounds:** Frequency bounds are set to $\pm 0.1\%$ of the spectral guess. All amplitudes and phases are left free within the normalised domain. A Trust-Region Reflective (TRF) solver with a linear loss is used.

2. **Refining:** Bounds are tightened to within $\pm 2\%$ of **each** parameter returned by the previous run. The same TRF solver is used.

3. **Final:** Bounds are tightened further to $\pm 1\%$ of previous results, and the fit is now run on the full, unclipped data window. This final pass extracts the best-estimate parameters and the covariance matrix.

The core fitting model and the pass sequence are shown in Listing A.5. Some samples of the fits used for debugging the fitting code are shown in Figure A.4.

```
def fit_triple_sine_linear(t, x, freq_guess, bounds=None):
    # -- Normalise --
    x_mean, x_std   = np.mean(x), np.std(x)
    t_min, t_range  = np.min(t), np.max(t) - np.min(t)
    t_norm = (t - t_min) / t_range
    x_norm = (x - x_mean) / x_std
    wg = freq_guess * t_range          # scaled frequency guesses

    def model(t, A1,w1,p1, A2,w2,p2, A3,w3,p3, off, slope):
        return (A1*np.sin(w1*t+p1) + A2*np.sin(w2*t+p2)
               + A3*np.sin(w3*t+p3) + off + slope*t)

    dw = 0.001   # initial frequency window (+/- 0.1%)
    lo = [1e-6, wg[0]*(1-dw), -np.pi,  1e-6, wg[1]*(1-dw), -np.pi,
          1e-6, wg[2]*(1-dw), -np.pi,  -3, -np.inf]
    hi = [5.0,  wg[0]*(1+dw),  np.pi,  5.0,  wg[1]*(1+dw),  np.pi,
          5.0,  wg[2]*(1+dw),  np.pi,   3,  np.inf]

    # Wide
    p1, _ = curve_fit(model, t_norm[1:-1], x_norm[1:-1],
                      bounds=[lo, hi], method='trf', maxfev=1500000)

    # tighten to +/-2% of Prev. solution
    lo2, hi2 = _tighten_bounds(p1, pct=0.02, Alo=1e-6)
    p2, _ = curve_fit(model, t_norm[1:-1], x_norm[1:-1],
                      bounds=[lo2, hi2], method='trf', maxfev=1500000)

    # tighten to +/-1%, full data
    lo3, hi3 = _tighten_bounds(p2, pct=0.01, Alo=1e-6)
    popt, pcov = curve_fit(model, t_norm, x_norm,
                           bounds=[lo3, hi3], method='trf', maxfev=1500000)
```

```

    # -- Unscale to original units --
    popt_out = _unscale_triple(popt, x_mean, x_std, t_range)
    return popt_out, _unscale_cov(pcov, x_std, t_range)
```

Listing A.5: Three-pass triple-sine fit with iterative bound tightening (normalised domain, simplified for clarity).

Once ω_1 is determined from the post-drop fit, the pre-drop segment is fitted with a single-mode model at that fixed frequency:

$$x_1(t) = A'_1 \sin(\omega_1 t + \varphi'_1) + c'_0 + c'_1 t, \tag{A.6}$$

with only $\{A'_1, \varphi'_1, c'_0, c'_1\}$ as free parameters. This is done because the cooled pre-drop signal is weak, the particle is in the cooled steady state, and (nearly) no radial mode signal is present. The code for this step is shown in Listing A.6.

```python
def fit_sine_fixed_freq(t, x, freq_fixed=90e3*2*np.pi,
                        bounds=([0.001, -np.pi, -np.inf, -np.inf],
                                [np.inf,  np.pi,  np.inf,  np.inf])):
    # model: A*sin(freq_fixed*t + phi) + offset + slope*t
    def model(t, A, phi, offset, slope):
        return A * np.sin(freq_fixed * t + phi) + offset + slope * t

    A_guess     = np.std(x) * np.sqrt(2)
    p0 = [A_guess, np.pi, np.mean(x), 0.0]

    try:
        popt, pcov = curve_fit(model, t, x, p0=p0, bounds=bounds,
                               method='trf', loss='soft_l1',
                               f_scale=1, maxfev=15000)
        return popt, pcov   # (A, phi, offset, slope)
    except RuntimeError:
        return None, None
```

Listing A.6: Single-mode fixed-frequency fit for the pre-drop segment.

Rather than using the raw signal variance, which mixes all three mechanical modes plus slope, the mode variance is computed analytically from the fitted axial amplitude,

$$\sigma_k^2 = \frac{A_k^2}{2}, \tag{A.7}$$

which is the time-averaged mean-square displacement of a single sinusoidal mode. In practice, the variance is estimated numerically by evaluating $A_k \sin(\omega_k t + \varphi_k)$ over a synthetic time vector spanning many full oscillation periods and computing its variance, which converges to $A_k^2/2$ irrespective of the phase φ_k.
These two amplitudes are then squared, averaged over the ensemble of shots, and divided to form the relative mean energy $\langle E \rangle / \langle E_0 \rangle$, as discussed in the caption and paragraph of Figure 4.21.

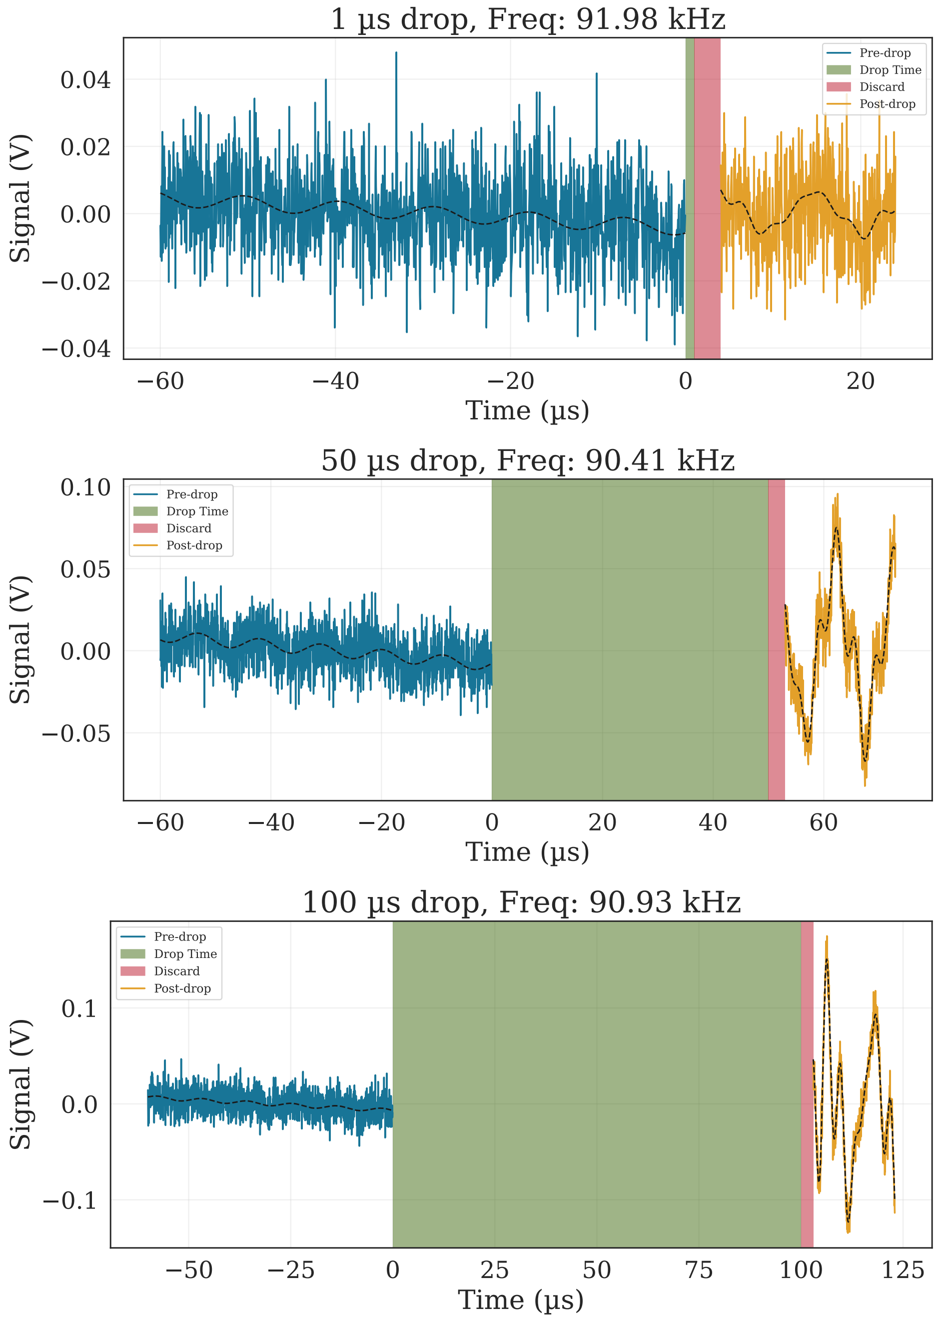

Figure A.4.: **Sine Fits for mean energy analysis.** These timetraces are representative random samples at different drop times τ for checking whether the fitting procedure is working. The blue line is the pre drop data, while the orange line is the post drop data, both filtered. The red shaded area is discarded data to account for detector dead time. The green shaded area is the drop time, while the violet line in the background is the unfiltered, raw post drop data for comparison. In dashed black, the fits pre and post drop are shown. The *Freq:* in the title is the extracted frequency from the post drop fit, which is then used as a fixed parameter for the pre drop fit.

A.2.6. Voltage Scan: PSD Integration

Used in Section 4.3. Alternatively, energy is calculated by integrating the Power Spectral Density (PSD) of the signal. This approach is particularly advantageous when the signal contains multiple frequency components or significant noise, as it allows isolating the energy within a specific bandwidth. It should therefore be more robust in the low τ regime, where the signal is mostly noise due to the cooled state, whereas the sine fitting struggles to find the correct frequency and amplitude.
The algorithm first computes the periodogram using a Hann window to reduce spectral leakage. It then identifies the dominant peak corresponding to the particle's motion using a peak-finding algorithm. Finally, the spectral density $S_{xx}(f)$ is integrated within a defined bandwidth BW centred around the detected peak frequency f_c ($BW = 30\,\mathrm{kHz}$ for the axial z mode, $3\,\mathrm{kHz}$ for transverse x/y modes).

$$V = \int_{f_c-BW}^{f_c+BW} S_{xx}(f)\, df \tag{A.8}$$

In the discrete implementation, this integration yields the signal variance (mean square amplitude) directly.

```
# Calculate Periodogram with Hann window
f, Pxx = signal.periodogram(x, fs, window='hann', nfft=nfft)

# Find peak frequency
peaks, _ = signal.find_peaks(np.log10(Pxx), ...)
center_freq = f[peaks[0]]

# Define integration mask and integrate
# Bandwidth is defined as +/- bw (e.g., +/- 30kHz)
mask = (f >= center_freq - bw) & (f <= center_freq + bw)
variance = np.trapezoid(Pxx[mask], f[mask])
```

Listing A.7: PSD Integration Logic

A.2.7. Squeezing pulse scan analysis details

For the squeezing scan in Figure 4.26, the analysis is similar to the recompression, but as mentioned in the main text, the large excitation in the radial modes means the model we use to motivate the separate analysis breaks down (see Section 2.3). Due to the different frequencies in the modes, the pulse times leading to the best attempt at squeezing for the z mode are also those that very strongly expand the radial modes. By using the first pulse time as a reference, we exclude, for all subsequent runs, all data points with a radial amplitude larger than `medScale * MAD`, so a scale factor (we found medScale = 3 to work well) times the median absolute deviation (MAD) of the radial amplitudes. Additionally, we post-select on frequencies within $0.9 \times 8,\mathrm{kHz}$ of the initial run.
We see in Figure A.5 the number of points that are kept for the squeezing analysis, following the behaviour as expected for radial modes that get progressively more excited. Since we post-select on a different measure than the one we are interested in, we can be

sure not to bias the result, as with this post-selection we would also remove those points where cross-coupling would show up.

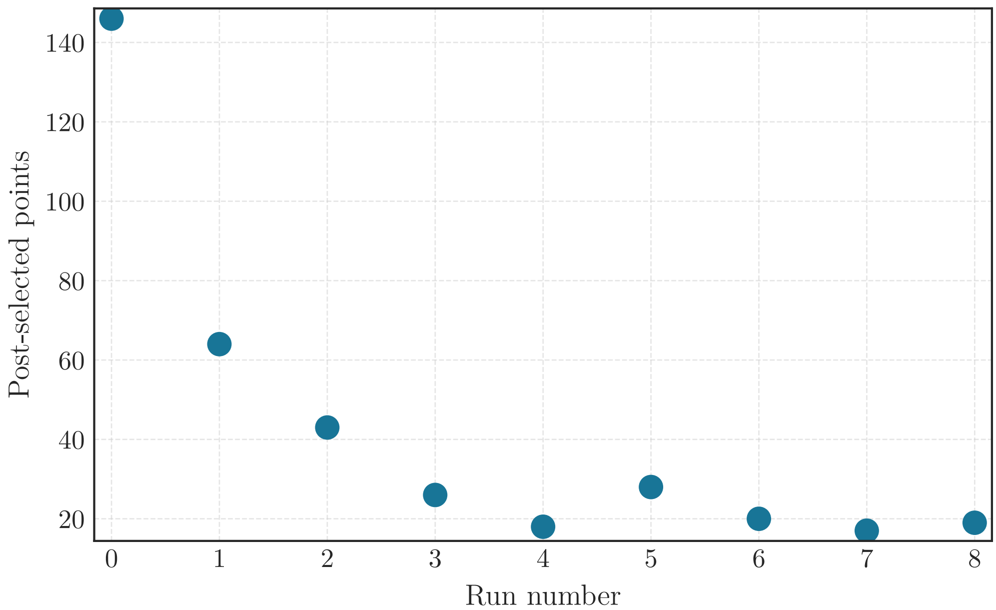

Figure A.5.: **Kept points after post-selection in squeezing analysis.** The number of points that are kept for the squeezing analysis after post-selecting on the radial amplitude. The first pulse time is used as a reference to set the threshold for post-selection, and we see that as the pulse time increases, the number of kept points decreases, following the expected behaviour for continuously increasing radial excitation. The low number of points for the optimal pulse time

A.2.8. Bootstrapping

Developed by Bradley Efron in the late 1900s [67], bootstrapping is used to gain information about the distribution of an estimator by resampling. In terms that are potentially less accurate but hopefully more intuitive, let us phrase the problem: We have a population N, of which we have measured a parameter n times (our sample). We are now interested in the mean of this parameter. Of course, we use the n samples we took and take the mean of all of them, giving us the best possible value for the mean, but now we have used all our data and know nothing about how large the uncertainty on this mean is. This is the basic problem bootstrapping (or jackknifing, the method bootstrapping evolved out of) is here to solve.
As is usually the case, we assume a few things about the set of observations: they need to be from an independent and identically distributed population (IID). This is assumed to be true for our observations, as independence is given by the fact that we let the particle

re-initialise (non-Markovian would be the word statistical physics uses) and identically distributed means that we sample from the same distribution (the process that makes our distribution is the same for all realisations). An important confusion that needs to be cleared up: the distribution for *identically* distributed is not necessarily an *equal* distribution.
Now bootstrapping does the following, taking the mean as an example (other properties of a sample can also be treated the same way): while the sample was previously everything we had, we now ignore the population and treat the sample as if it were the new population. We take the data, and sample again with replacement (as many subsamples as possible, we typically use 300, as we do not rely heavily on the error bars and computational constraints), generating our subsamples. We have now solved our problem: by looking at the distribution of the mean of our subsamples, we can put error bars on the mean calculated from all the data, the original sample.

A.3. Lyapunov Covariance Propagation

The squeezing and recompression analysis scripts simulate the expected covariance matrix evolution using a numerical implementation of the Lyapunov differential equation introduced in Equation (2.18). As pointed out in the main text, as the positive results do not depend on this theory, this section documents that implementation as it is not just a direct solving of the equation.

Vectorisation

The covariance matrix $\mathbf{\Sigma}$ is 2×2 and symmetric, so it has only three independent entries. Rather than propagating the full matrix, we represent it as a 3-vector

$$\mathbf{p} = \begin{pmatrix} \Sigma_{xx} \\ \Sigma_{xp} \\ \Sigma_{pp} \end{pmatrix}, \tag{A.9}$$

and the Lyapunov equation $\dot{\mathbf{\Sigma}} = \mathbf{A\Sigma} + \mathbf{\Sigma A}^{\mathrm{T}} + \mathbf{D}$ becomes the linear, inhomogeneous system

$$\dot{\mathbf{p}} = \mathbf{Mp} + \mathbf{q}, \tag{A.10}$$

where $\mathbf{M}$ is a 3×3 matrix and $\mathbf{q}$ a 3-vector obtained from $\mathbf{A}$ and $\mathbf{D}$ by expanding the products entry by entry:

$$\mathbf{M} = \begin{pmatrix} 2A_{11} & 2A_{12} & 0 \\ A_{21} & A_{11} + A_{22} & A_{12} \\ 0 & 2A_{21} & 2A_{22} \end{pmatrix}, \qquad \mathbf{q} = \begin{pmatrix} D_{xx} \\ D_{xp} \\ D_{pp} \end{pmatrix}. \tag{A.11}$$

In code, this vectorisation is:

```
def get_lyapunov_matrix(A):
    return np.array([
        [2*A[0,0],          2*A[0,1],          0        ],
        [   A[1,0], A[0,0]+A[1,1],        A[0,1]      ],
        [        0,         2*A[1,0],     2*A[1,1]      ],
    ])
```

Listing A.8: Vectorisation of the drift matrix into the Lyapunov propagation matrix.

Propagation via the matrix exponential

The inhomogeneous ODE in Equation (A.10) is solved exactly by augmenting the system by one dimension, a process similar to variation of parameters that includes the inhomogeneous term in the system matrix, resulting in a homogeneous system:

$$\frac{\mathrm{d}}{\mathrm{d}t}\underbrace{\begin{pmatrix}\mathbf{p}\\ 1\end{pmatrix}}_{\mathbf{p}(t)} = \underbrace{\begin{pmatrix}\mathbf{M} & \mathbf{q}\\ \mathbf{0}^{\mathrm{T}} & 0\end{pmatrix}}_{\mathbf{A}_{\mathrm{aug}}}\begin{pmatrix}\mathbf{p}\\ 1\end{pmatrix}. \tag{A.12}$$

This is now a homogeneous linear system $\frac{\mathrm{d}}{\mathrm{d}t}\mathbf{p}(t) = \mathbf{A}_{\mathrm{aug}}\mathbf{p}(t)$ whose solution at time t is $e^{\mathbf{A}_{\mathrm{aug}}t}$ applied to the initial state, which `scipy.linalg.expm` computes directly [68]:

```
def propagate_cov(P0, M, q, t):
    p0  = np.array([P0[0,0], P0[0,1], P0[1,1]])
    aug = np.zeros((4, 4))
    aug[:3, :3] = M
    aug[:3,  3] = q
    p_t = (expm(aug * t) @ np.append(p0, 1.0))[:3]
    return np.array([[p_t[0], p_t[1]], [p_t[1], p_t[2]]])
```

Listing A.9: Covariance propagation via matrix exponential using SciPy expm.

Protocol sequence

The full propagation (free evolution, rotation pulse in the trap, second free evolution) is now just three calls to that propagation function, plugging in the previous result:

```
P_tau1  = propagate_cov(P0,      M_free, q_free, tau_1)  # free fall
P_pulse = propagate_cov(P_tau1,  M_trap, q_trap, t_p  )  # harmonic pulse
P_tau2  = propagate_cov(P_pulse, M_free, q_free, tau_2)  # free fall
```

where the drift and diffusion matrices are built as:

```
A_free = np.array([[0,           1/m ],
                   [0,             0 ]])      # gamma_gas = 0 due to
    negligible effect on the timescales considered

A_trap = np.array([[0,           1/m ],
                   [-m*Omega**2,   0 ]])

D_pp_gas = 4 * Gamma_gas * Omega**2 * m**2  # gas diffusion
```

```
D_pp_rec = 4 * Gamma_rec * Omega**2 * m**2  # photon recoil

Q_free = np.array([[0, 0], [0, D_pp_gas           ]])
Q_trap = np.array([[0, 0], [0, D_pp_gas + D_pp_rec ]])
```

A.4. Recompression over multiple minima

Following the same experimental procedure as described in Section 4.2, we can resolve multiple minima corresponding to additional integer 2π rotations in phase with respect to the rotation in Figure 4.4. These data are taken with worse initial cooling ($n \approx 40$) and with fewer repetitions (N=50). The free evolution time here is $\tau = 5\,\mu\text{s}$ before and after the pulse t'.

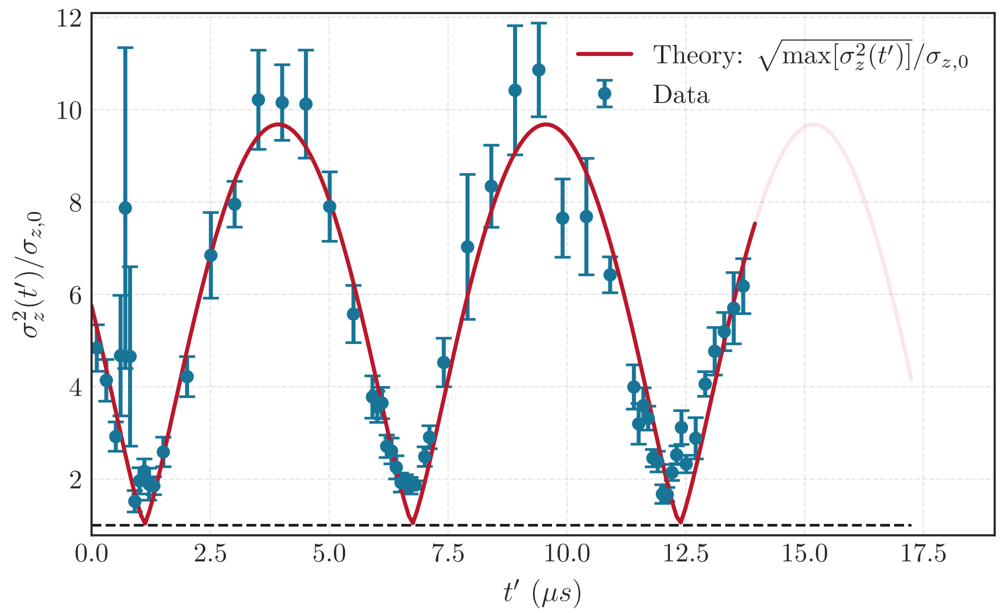

Figure A.6.: **Recompression over multiple minima:** The number of runs for each data point is $n = 50$, the other settings are the same as in Figure 4.5

A.5. Force Calibration

To convert applied voltages to forces, we perform a calibration as done in [12, Appendix C]. Having a calibrated position detection (as explained in Section A.1), we can already convert a homodyne voltage signal to a position signal. If we now drive the particle, far off resonance, with a known voltage, we have all the ingredients to extract the conversion factor C_{NV}. This factor links the force F_{drive} applied to the particle to the voltage V_{drive} we apply to the electrodes, such that $F_{\mathrm{drive}} = C_{\mathrm{NV}} V_{\mathrm{drive}}$.

It is important to drive off resonance, as natural particle motion would make this treatment more complicated. For our particle frequencies, we chose to drive at 160 kHz and at 210 kHz. By then measuring many force and voltage pairs, we can fit a line with slope C_{NV} to the data. In our setup, this slope comes out to be $C_{\mathrm{NV}} = 2.31 \times 10^{-16}\,\mathrm{N\,V^{-1}}$ and $2.25 \times 10^{-16}\,\mathrm{N\,V^{-1}}$ for the two frequencies. By measuring at two frequencies, we would see a contribution of the Ω_z mode in our calibration, as this would lead to an asymmetry in slopes.

A.6. Used Equipment

This list provides a quick reference for the equipment used. While every piece of equipment is introduced in the main text, this list can be used to quickly look up specific model numbers and properties when reading through a chapter. Mentions in the main text are kept more general to avoid breaking the flow.

Component	**Model / Details**
Digital Oscilloscope	Picoscope 5000Series
DC Voltage Supplies	Stanford DC205
High Voltage Power Supply	RS-A3305P
RF Switches	Mini-Circuits ZASWA-2-50DRA+
Pulse Generator	Keysight 33500B
Bias Tee	In-house built (Design: Gregor Maier)
X,Y DC to RF Amplifiers	Falco Systems WMA-100
Balance Detectors	ThorLabs PDB425C
Photodiode	ThorLabs PDA10CF-EC
PID Controller	In-house built
Laser	Coherent Mephisto 1064 nm
AOM Drive	Wieserlabs FlexDDS-NG Dual Channel Phase Coherent 400MHz RF
AOM Drive Amplifier	Mini-Circuits ZHL-03-5WF+
Red Pitaya	STEMLab 125-14
Lock-in Amplifier	Zurich Instruments HF2LI

Table A.1.: Equipment used in the experiment.

A.7. Optical Setup of the Split Detection

As outlined in Section 3.1.5, the forward detection is split into three paths, one for each mode. Here we sketch the beam path on top of a picture of the actual setup, to give a better idea of how the light is split and directed to the different detectors.

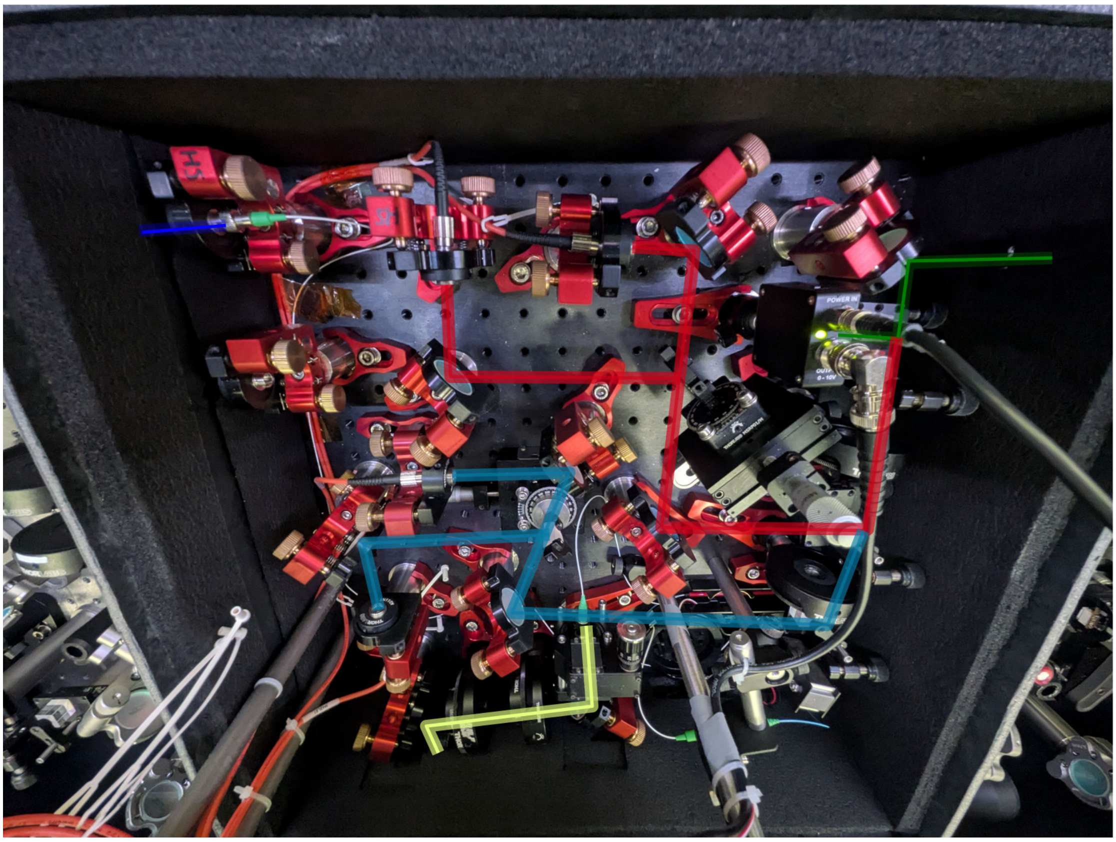

Figure A.7.: **Optical setup for the forward detections.** In their respective colours we have shown the x (light blue) and y (red) split detection paths. The light in each path is passed through a WP and PBS to split off a part, then reflected partly on a D-shaped mirror. In green we have marked the path of the forward z detection onto a single photodiode. From below we see the BW scattered light entering the detection box and being coupled into fibres to then form the homodyne detection. On the top left, we see the LO path, which was split off before the vacuum chamber and is also coupled into a fibre in the detection box.

List of Tables

Listings