University of Padova

Department of Physics "Galileo Galilei"

Master Thesis in Physics of Data

Trapped-ion quantum architectures:

RF electrodes optimization and effects of RF technical noise

Supervisors

Prof. Pietro Silvi

University of Padova

Master Candidate

Lorenzo Calandra Buonaura

Student ID

2107761

Co-supervisors

Dr. Philipp Schindler

University of Innsbruck

Marco Valentini

University of Innsbruck

Academic Year

2024-2025

χαλεπὰ τὰ καλά

— Plato

Abstract

Ion traps are among the most promising platforms for quantum computing, owing to the record gate fidelities and long coherence times demonstrated on such devices. Currently, most trapped-ion quantum processors are realized using Paul traps, where a combination of radio-frequency (RF) and direct-current (DC) voltages applied to the trap electrodes is used to confine and control ions in space. This technology can be scaled towards architectures capable of controlling hundreds or thousands of ions by exploiting standard microfabrication techniques to miniaturize the electrodes and integrate them on the surface of a chip, forming the so-called two-dimensional surface-electrode ion traps.

However, scaling up the number of ions controlled on a surface Paul trap while preserving high computational fidelity remains a major challenge. Efficient and independent control of many ions requires small ion–surface distances, which increases the susceptibility of the ions to noise originating from the trap electrodes. This technical noise can couple to the ions' motional degrees of freedom, inducing heating and decoherence that degrade gate performance. While the effects of DC voltage noise on the ion motion have been extensively studied in the literature, the impact of technical noise on the RF drive remains comparatively less understood and explored.

The purpose of this thesis is to develop a theoretical framework that characterizes how RF technical noise affects the secular motion of trapped ions, leading to motional heating. In the first part of the thesis, a theoretical model describing the effects of RF phase and amplitude noise in terms of ion heating is developed, highlighting the role of the noise spectral density and the trap geometry in determining the heating rates experienced by the ions. In the second part, the developed model is applied to two trapped-ion quantum computing architectures: the Quantum Charge-Coupled Device (QCCD) architecture and the Quantum Spring Array (QSA) architecture. The model is first applied to the electrode layout of an X-junction trap, a fundamental zone in the QCCD architecture where two linear traps intersect. It is shown how the optimization of the RF electrode edges reduces heating and other detrimental effects experienced by ions during shuttling across the junction. The model is then applied to RF transport, a fundamental ion-shuttling mechanism used to connect different zones in the QSA architecture. Heating rates are simulated along the shuttling path using the developed models and are compared with experimental data, showing reasonable agreement between the two.

Abstract

Le trappole per ioni sono tra le piattaforme più promettenti per la computazione quantistica, grazie alle elevate fedeltà nelle porte logiche e ai lunghi tempi di coerenza dimostrati su questi dispositivi. Al momento, la maggior parte dei processori quantistici a ioni intrappolati sono realizzati utilizzando trappole di Paul, dove la combinazione di voltaggi a radiofrequenza (RF) e a corrente continua (DC) applicati sugli elettrodi della trappola è utilizzata per confinare e controllare gli ioni nello spazio. Questa tecnologia può essere ampliata verso architetture capaci di controllare centinaia o migliaia di ioni sfruttando tecniche di microfabbricazione standard per miniaturizzare gli elettrodi e integrarli sulla superfice di un chip, creando le cosiddette trappole per ioni bidimensionali con elettrodi superficiali.

Tuttavia, aumentare il numero di ioni controllati su una trappola di Paul superficiale preservando un'elevata fedeltà computazionale rimane una sfida importante. Il controllo efficiente e indipendente di molti ioni richiede piccole distanze tra gli ioni e la superficie, il che aumenta la suscettibilità degli ioni al rumore proveniente dagli elettrodi della trappola. Questo rumore tecnico può accoppiarsi ai gradi di libertà di movimento degli ioni, inducendo riscaldamento e decoerenza che degradano le prestazioni delle porte logiche. Mentre gli effetti del rumore della corrente continua sul movimento degli ioni sono stati ampiamente studiati nella letteratura, l'impatto del rumore tecnico sul segnale a radiofrequenza rimane relativamente meno compreso ed esplorato.

Lo scopo di questa tesi è quello di sviluppare un quadro teorico che caratterizzi il modo in cui il rumore tecnico a radiofrequenza influisce sul moto secolare degli ioni intrappolati, causando riscaldamento cinetico. Nella prima parte della tesi viene sviluppato un modello teorico che descrive gli effetti del rumore RF di fase e ampiezza per quanto riguarda il riscaldamento degli ioni, evidenziando il ruolo della densità spettrale del rumore e della geometria della trappola nel determinare i tassi di riscaldamento a cui sono soggetti gli ioni. Nella seconda parte, il modello sviluppato viene applicato a due architetture di calcolo quantistico a ioni intrappolati: l'architettura Quantum Charge-Coupled Device (QCCD) e l'architettura Quantum Spring Array (QSA). Il modello viene prima applicato al layout degli elettrodi di una trappola a giunzione X, una zona fondamentale nell'architettura QCCD dove due trappole lineari si intersecano. Si mostra come l'ottimizzazione dei bordi dell'elettrodo RF riduce il riscaldamento e altri effetti dannosi subiti dagli ioni durante il movimento attraverso la giunzione. Il modello viene quindi applicato al trasporto RF, un meccanismo fondamentale per il trasporto di ioni utilizzato per connetter diverse zone nell'architettura QSA. I tassi di riscaldamento vengono simulati lungo il percorso di trasporto utilizzando i modelli sviluppati e vengono confrontati con i dati sperimentali, mostrando una ragionevole concordanza tra i due.

Contents

List of Figures

List of Tables

List of Acronyms

COM . . . . . center of mass

DC . . . . . direct-current

PS . . . . . pseudopotential

PSD . . . . . power spectral density

QCCD . . . . . Quantum Charge-Coupled Device

QSA . . . . . Quantum Spring Array

RF . . . . . radio-frequency

TS-VIAs . . . . . through-substrate vertical interconnect accesses

1

Introduction

The origin of modern computer science is often attributed to the great mathematician Alan Turing, who, in his remarkable 1936 paper, introduced an abstract model of computation, now known as the *Turing machine* [1]. This theoretical framework laid the foundation of classical computer science, providing a formal description of what it means to compute, and established the conceptual basis for *programmable machines*. Since then, classical computers have evolved dramatically, becoming faster, more compact, and more energy-efficient, and are now ubiquitous in our daily lives. However, in 1982, Richard Feynman pointed out that there appeared to be fundamental difficulties in simulating quantum mechanical systems using classical computers [2]. He suggested that building computers based on the principles of quantum mechanics would allow us to avoid those limitations, as nature itself is inherently quantum. Around the same time, Paul Benioff independently proposed a quantum mechanical model of the Turing machine [3], laying the groundwork for what would become the field of quantum computation.

From classical to quantum computation

Quantum computing represents a revolutionary shift in the way information can be processed. Unlike classical computers, whose basic unit of information is the *bit*, which can take on of two values, 0 or 1, quantum computers employ quantum two-level systems, or *qubits*, which can exist in a *superposition* of the quantum states $|0\rangle$ and $|1\rangle$. Moreover, multiple qubits can exhibit *entanglement*, a fundamental quantum phenomenon in which their states become correlated in ways that have no classical counterpart. This correlation allows the state of each qubit to depend on the others, enabling the representation of complex and nonlocal information structures. These uniquely quantum features can be used as practical computational tools, allowing some problems to be

solved more efficiently than on classical computers.

Given this radical departure from the classical paradigm, it is natural to ask what advantages quantum information processing can offer over the well-established technology of classical computation. Despite their ubiquity and continual advancement, classical computers are approaching fundamental physical limits, such as those imposed by Landauer's principle [4], which sets a lower bound on the energy required to erase one bit of information. Moreover, even if hardware improvements continue, there remain classes of problems for which no efficient classical algorithms are known. Quantum computers, by exploiting superposition and entanglement, can in principle solve some of these problems more efficiently. Prominent examples include Shor's algorithm for integer factorization [5] and Grover's search algorithm [6], which respectively offer exponential and quadratic speed-ups compared to their best-known classical counterparts. The potential of quantum computation, however, extends far beyond a handful of known algorithms. As originally envisioned by Feynman [2], a quantum processor can efficiently simulate the dynamics of another quantum system, a task that quickly becomes impossible for classical computers. For instance, storing the full quantum state of only 80 spin-$\frac{1}{2}$ particles would require a classical memory of 2^{80} complex numbers, corresponding to approximately 5×10^{12} TB, more than the total of what was globally stored in 2007 [7]. Problems of this kind appear in many areas of physics and chemistry, such as modelling quantum phase transitions, high-T_c superconductivity, or molecular reactions relevant to nitrogen and carbon dioxide fixation [7, 8, 9].

Quantum computing platforms

Over the past decades, significant progress has been made towards realizing physical platforms capable of performing quantum information processing. A wide range of experimental implementations are currently being investigated, including superconducting circuits [10], neutral atoms [11], photons [12], nuclear spin ensembles [13], defect centers in solid-state systems [14], and trapped ions [15]. Each of these platforms fulfils the fundamental requirements for quantum computation outlined by DiVincenzo [16].

Every architecture is more suitable for specific tasks depending on its physical properties and operational strengths. Superconducting qubits, for instance, are controlled using conventional microwave electronics, which allows for fast gate operations and straightforward integration into scalable architectures. Neutral atoms, on the other hand, offer exceptional isolation from the environment when idle, enabling remarkably long coherence times. Photonic qubits, being inherently mobile, are ideal carriers

of quantum information and thus are suitable for quantum communication protocols. Nuclear spin systems are extremely stable and highly sensitive, making them excellent candidates for quantum sensing, while defect centers in solids combine optical addressability with the robustness of solid-state platforms.

Finally, *trapped ions* represent one of the most mature and versatile technologies for both quantum computation and simulation. Qubits are encoded in the internal electronic states of atomic ions and manipulated using laser or microwave radiation [15, 17]. Being charged particles, ions can be confined for long periods of time within electromagnetic potentials, remaining well isolated from the environment and preserving quantum coherence. For example, in a linear Paul trap [18], a combination of static (DC) and oscillating (RF) electric fields, produced by a specific arrangement of electrodes, creates a time-averaged potential that traps ions along a line while allowing precise control of their motion. The mutual Coulomb repulsion between ions leads to ordered structures that enable individual optical addressing and provide intrinsic qubit connectivity. State-of-the-art trapped-ion quantum processors based on macroscopic linear Paul traps can store and manipulate up to a few tens of ion qubits with high fidelity [19, 20, 21, 22], allowing the demonstration of fundamental quantum algorithms, quantum error correction protocols, and both analog and digital quantum simulations [23, 24, 25, 26].

Quantum information processing with trapped ions

In trapped-ion quantum computers, quantum information is encoded in the long-lived internal electronic states of atomic ions, while their collective motion provides a natural mechanism for mediating interactions between qubits [17, 27]. When multiple ions are confined in the same trapping potential, their mutual Coulomb repulsion arranges them into a linear crystal with well-defined equilibrium positions. All-to-all connectivity between the ion qubits in the chain is enabled by their shared vibrational motion. Among these collective modes, the center-of-mass (COM) mode is widely used for information processing, as in this mode all ions oscillate in phase and couple uniformly to external control fields. The information stored in the internal electronic states can be manipulated, for instance, using laser-driven interactions, in which the absorption or emission of photons transfers momentum to the ions and couples their internal states to the collective motion [15, 28]. By appropriately tuning these interactions, it is possible to implement high-fidelity single-qubit operations as well as multi-qubit entangling gates that act across the entire ion chain. Well-known entangling gate schemes, such as the Mølmer–Sørensen [29, 30] and Cirac–Zoller [31] gates, exploit this mechanism

by using the shared vibrational motion as a *quantum bus* to mediate interactions and enable universal quantum computation.

Challenges towards scalable trapped-ion quantum computers

Trapped-ion platforms currently hold some of the highest reported fidelities for single- and two-qubit gates among all quantum computing architectures (for instance, single-qubit fidelities above 99.999% [32] and two-qubit fidelities reaching 99.99% or better [33]). Still, a major challenge for trapped-ion quantum computers is to scale up the number of qubits controlled by such devices while simultaneously maintaining the performance of gate operations at high fidelities. In linear Paul traps, for example, as the ion chain grows longer, information processing at high fidelity becomes increasingly difficult due to limitations in optical access, motional-mode crowding, and geometric constraints [34, 35]. To address these issues, *surface-electrode ion traps* have emerged as a promising platform for realizing large-scale quantum processors [17, 20, 36, 37, 38, 39]. These traps are produced through modern microfabrication techniques, which enable the realization of complex trap architectures suitable for scalable quantum information processing. However, achieving scalability through miniaturization introduces new technical challenges. As electrode dimensions are reduced to allow for denser architectures, the local control of each ion qubit demands smaller ion-surface distances. A smaller separation between the ions and the trap makes ion qubits more susceptible to electric field noise coming from the trap surface, and this noise can lead to heating of the ion motion, which in turn limits gate fidelities and the overall system performance [40, 41]. While *surface-induced electric field noise* and technical noise on the trap DC electrodes have been extensively investigated as sources of motional heating in microfabricated surface traps [40, 42, 43, 44], the contribution of *RF technical noise* remains comparatively unexplored. The characterization of the relation between RF technical noise and ion-heating is important to improve the fidelity of quantum operations in surface-electrode traps and constitutes the main focus of this thesis.

Outline of the thesis

The goal of this work is to develop a theoretical model of the effects of technical RF noise on trapped ions, with particular emphasis on its contribution to motional heating. The thesis is structured as follows. Chapter 2 introduces the fundamental theoretical concepts underlying trapped-ion quantum processors and surface-trap architectures. It discusses the confinement of charged particles through the combined action of DC and

RF fields in a Paul trap and the scaling of such systems to larger ion numbers. In addition, it describes the relevant electronic level structure of the $^{40}Ca^{+}$ ion, provides an overview of the fundamentals of ion–light interactions, and introduces the concept of motional heating. In Chapter 3, the theoretical framework describing the heating induced by technical RF noise is developed, including analytical models for both amplitude and phase fluctuations. Chapter 4 and Chapter 5 then investigate two specific architectures in which RF technical noise plays a significant role. Among the various trapped-ion quantum computing architectures, this work focuses on the *Quantum Charge-Coupled Device* (QCCD) and the *Quantum Spring Array* (QSA). In the QCCD architecture, ion qubits are interconnected through a combination of shuttling, crystal rotation, and splitting or merging operations, enabling reconfigurable connectivity within a segmented trap [34]. In contrast, the QSA architecture consists of a two-dimensional array of parallel linear traps, each providing multiple independently controllable trapping sites, where ions can be transported and coupled along both axial and radial directions [45]. For the QCCD architecture, the analysis focuses on X-junctions, where geometric optimization of the RF electrode shapes can be employed to minimize heating effects during ion transport. For the QSA architecture, the study concentrates on RF-transport mechanisms, applying the previously developed heating models to simulate ion heating within the trap. The simulated results are then compared with experimental measurements to assess and validate the theoretical framework. Finally, Chapter 6 presents a brief summary of the main results obtained throughout this thesis, discusses their broader implications for the development of scalable trapped-ion quantum processors, and outlines possible directions for future research. In addition, Appendix A provides a supplementary analysis of the X-junction. Specifically, Section A.1 reports the optimization results for different ion–surface distances, Section A.2 examines the DC electrodes layout and the associated transport waveforms during shuttling, and Section A.3 discusses strategies for optimizing the routing of the wires connecting the electrodes to the signal supplies. The computational tools used for the simulations supporting this work are provided in Appendix B.

2

Ion traps as a platform for quantum information processing

In this chapter, we present the theoretical background required to understand the rest of the thesis, where the effects of technical RF noise on trapped ions are theoretically (Chapter 3) and quantitively (Chapter 4 and Chapter 5) analyzed. The purpose of this chapter is to establish the physical framework necessary to describe ion trapping, thereby clarifying why motional heating arises in ion traps and why it represents a crucial factor in development of high fidelity trapped ion quantum computers.

The first section explains how particle confinement is achieved in a Paul trap, introducing key concepts such as the pseudopotential and the secular frequencies. The second section focuses on surface-electrode traps and outlines the theoretical principles underlying the numerical simulations carried out in this work. Finally, the last section provides an overview of how qubits can be realized using trapped ions, which is essential for understanding the experimental setup described in Subsection 5.4.1.

2.1 Paul traps

Ions are electrically charged particles and therefore experience forces in the presence of electric fields: this fundamental property allows them to be confined and manipulated with high precision using suitable electromagnetic potentials. The ability to trap ions for extended periods of time while maintaining excellent control over their motion is one of the key enabling techniques of modern atomic physics and quantum information experiments. Among the various trapping approaches, *Paul traps* [18], also known as *radio-frequency* (RF) traps, are among the most widely used. They achieve three-dimensional confinement by combining static (DC) and oscillating (RF) electric fields,

which together create an effective potential capable of trapping ions at well-defined equilibrium positions in space.

2.1.1 Radial confinement

To confine a charged particle at a fixed point in space, the easiest way would ideally be to generate a static electric field that provides restoring forces along all three spatial dimensions. A natural mathematical form for such confining potential is a quadratic function:

$$V_{\mathrm{DC}}\left(\vec{r}\right) = \frac{1}{2}\left(\alpha_x r_x^2 + \alpha_y r_y^2 + \alpha_z r_z^2\right), \tag{2.1}$$

where α_k are the curvature coefficients of the electric potential along the orthogonal axes $k \in \{x, y, z\}$. Full confinement would require all $\alpha_k > 0$, ensuring that the resulting forces are directed toward the trap center in every dimension. However, *Earnshaw's theorem* states that such a potential cannot exist for purely static electric fields [46]. In free space, where no charge density is present, the potential must always satisfy Laplace's equation:

$$\nabla^2 V_{\mathrm{DC}} = 0, \tag{2.2}$$

which, when applied to the potential in Equation 2.1, leads to the constraint:

$$\sum_k \alpha_k = 0 \tag{2.3}$$

implying that at least one curvature coefficient must be negative, corresponding to an anti-confining direction along which the ion is repelled, proving the theorem.

Paul traps overcome this limitation by superimposing a time-varying RF field onto the static one. If this potential is harmonic and centered at the same point, it can be expressed as:

$$V_{\mathrm{RF}}\left(\vec{r}, t\right) = \frac{1}{2}\left(\beta_x r_x^2 + \beta_y r_y^2 + \beta_z r_z^2\right)\cos(\Omega_{\mathrm{RF}} t), \tag{2.4}$$

where β_k are the curvature coefficients associated with the RF field (which must independently satisfy Equation 2.2) and Ω_{RF} is the driving frequency. The combination of the static and oscillating fields enables stable confinement of charged particles, and reads as follows:

$$\begin{aligned} V\left(\vec{r}, t\right) &= V_{\mathrm{DC}}\left(\vec{r}\right) + V_{\mathrm{RF}}\left(\vec{r}, t\right) = \\ &= \frac{1}{2}\left(\alpha_x r_x^2 + \alpha_y r_y^2 + \alpha_z r_z^2\right) + \frac{1}{2}\left(\beta_x r_x^2 + \beta_y r_y^2 + \beta_z r_z^2\right)\cos(\Omega_{\mathrm{RF}} t). \end{aligned} \tag{2.5}$$

Among the different geometries that realize such a potential, the *linear Paul trap* is the most widely used configuration in trapped-ion experiments, due to its capability of confining long chains of ions for quantum information processing [27, 28]. In this design, four elongated hyperbolic-shaped electrodes are arranged symmetrically around the z-axis, generating the oscillating quadrupole field in the radial (xy) plane. The radial section of such an electrode arrangement is depicted in panel (a) of Figure 2.1.

This geometry exhibits cylindrical symmetry around the trap axis, allowing one to impose the following conditions:

$$\begin{cases} \alpha_x = -\alpha_y = \alpha, & \alpha_z = 0, \\ \beta_x = -\beta_y = \beta, & \beta_z = 0, \end{cases} \tag{2.6}$$

which satisfy Laplace's equation and reflect the fact that a restoring force in one radial direction necessarily produces an anti-confining force in the orthogonal one. If we now apply a RF and a DC voltage to the same opposing electrodes, the application of these symmetry conditions results in the following expression for the trap potential in the radial plane:

$$V(x, y, t) = \frac{1}{2}\left[\alpha(x^2 - y^2) + \beta(x^2 - y^2)\cos(\Omega_{\mathrm{RF}} t)\right]. \tag{2.7}$$

In this configuration, the RF quadrupole term dynamically stabilizes the ion motion along x and y, while the DC term provides additional control over the potential landscape. Importantly, the instantaneous RF field vanishes along the trap axis ($x, y = 0$),

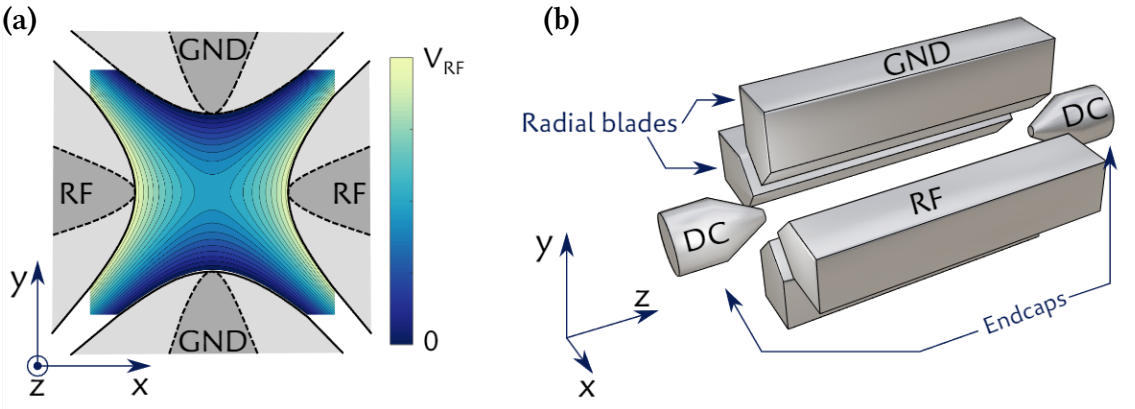

FIGURE 2.1: Overview of linear Paul trap geometry. (a) An applied voltage on one of two pairs of hyperbolic electrodes produces a harmonic saddle-potential, indicated by the contour lines. This potential becomes effectively confining when an oscillating voltage (RF) is applied. In reality, electrodes are not hyperbolic in shape, but more blade-like, as indicated by the dashed lines. (b) A positive voltage on endcaps placed on opposing ends of the z-axis provides axial confinement. Image courtesy of Martin van Mourik, adapted from his PhD thesis [47].

defining the so-called *RF null*. For the correct choice of RF and DC voltages, the ion remains confined within this null, and its motion can be described as small oscillations around the equilibrium position. The next section discusses the physical and experimental conditions required for the ion to remain confined within the trap.

2.1.2 Stable trajectories in the radial plane

The motion of an ion with charge Q and mass m in the radial plane of a linear Paul trap is governed by the combined RF and DC quadrupole potential introduced in Equation 2.7. This motion can be described using two uncoupled differential equations:

$$\left[m\frac{\mathrm{d}^2}{\mathrm{d}t^2} + Q\left(-\alpha + \beta\cos\left(\Omega_{\mathrm{RF}}t\right)\right)\right]x = 0, \tag{2.8a}$$

$$\left[m\frac{\mathrm{d}^2}{\mathrm{d}t^2} + Q\left(-\alpha + \beta\cos\left(\Omega_{\mathrm{RF}}t\right)\right)\right]y = 0. \tag{2.8b}$$

The two radial equations of motion are mathematically equivalent to the *Mathieu equation* [48, 49]:

$$\left[\frac{\mathrm{d}^2}{\mathrm{d}\xi^2} + \left(a - 2q\cos\left(2\xi\right)\right)\right]u = 0, \tag{2.9}$$

through the substitutions:

$$\xi = \frac{\Omega_{\mathrm{RF}}}{2}t, \qquad a_x = a_y = -\frac{4Q\alpha}{m\Omega_{\mathrm{RF}}^2}, \qquad q_x = q_y = \frac{2Q\beta}{m\Omega_{\mathrm{RF}}^2}. \tag{2.10}$$

This correspondence enables the use of the Mathieu stability analysis to determine the regions in which the ion motion is dynamically confined in the radial plane. The general solution of Equation 2.9 remains bounded as $\xi \to \infty$ only if the parameters a and q lie within specific *stability regions* of the (a, q) plane [50, 49]. These regions appear as alternating tongue-shaped areas in the (a, q) diagram, symmetric under the transformation $q \to -q$, so that only $q > 0$ is typically considered.

In a linear Paul trap, the motion along the radial directions is constrained by the symmetry conditions expressed in Equation 2.6. Stable confinement is therefore possible only in the parameter ranges where the stability bands for both x and y overlap; outside this overlap, the ion experiences parametric instability in at least one direction, leading to loss of confinement. Panel (a) in Figure 2.2 illustrates the global structure of these stability bands for $a \in [-5, 5]$ and $q \in [0, 10]$, showing the overall pattern of stable and unstable regions.

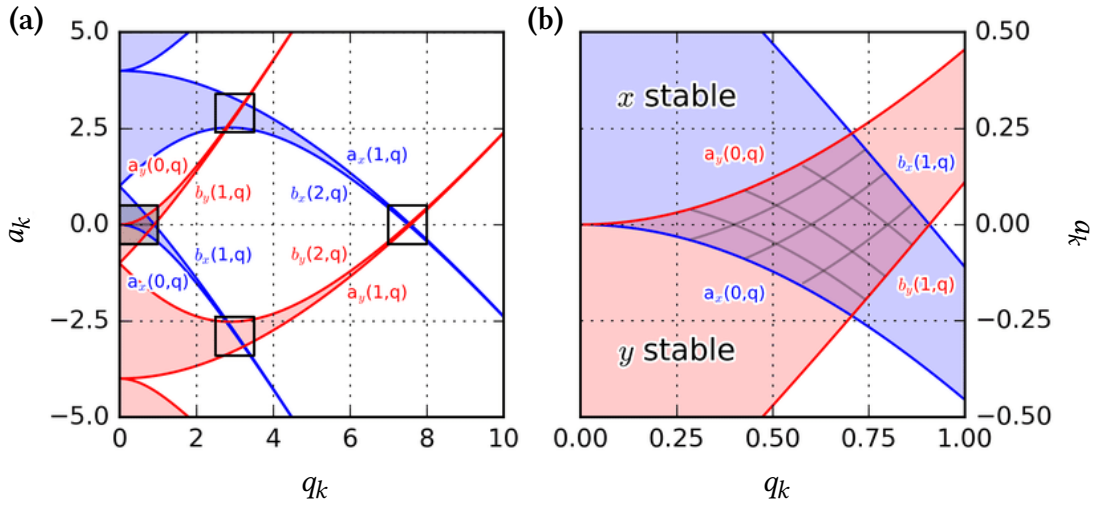

FIGURE 2.2: Stability diagram of the Mathieu equation. (a) Global view of the stability regions in the (a, q) plane for $a \in [-5, 5]$ and $q \in [0, 10]$. (b) Zoom on the lowest stability region for $a \in [-0.5, 0.5]$ and $q \in [0, 1]$, highlighting the overlap between the x and y stability bands where trapping is simultaneously stable along both radial directions [51].

For applications in quantum computing, ion traps are typically operated within the *lowest stability region*, corresponding to small stability parameters $|a_i|, q_i^2 \ll 1$ [28]. Panel (b) in Figure 2.2 provides a zoom into this region, relevant for practical trap operation. In the limit $|a_i| \ll q_i$, stable operation is approximately achieved for $0 \leq q_i \leq 0.908$. In real devices, deviations from the ideal quadrupole potential can introduce higher-order field terms, causing discrete instabilities at larger q values [52]. To ensure robust confinement in both radial directions, Paul traps are therefore operated with $q \lesssim 0.5$ in many designs [53], although in some cases even lower values ($q \approx 0.3$) are chosen to minimize micromotion and deviations from ideal quadrupole behavior [54, 55].

In the regime where $|a_i|, q_i^2 \ll 1$, the motion of the ion can be well approximated by the analytical solution [28]:

$$u_i(t) \approx A_i \cos(\omega_i t)\left[1 + \frac{q_i}{2}\cos(\Omega_{\mathrm{RF}} t)\right], \qquad i \in \{x, y\}, \tag{2.11}$$

where A_i is the oscillation amplitude and ω_i is the *secular frequency*, defined as:

$$\omega_i = \frac{\Omega_{\mathrm{RF}}}{2}\sqrt{a_i + \frac{q_i^2}{2}} \ll \Omega_{\mathrm{RF}}\,, \tag{2.12}$$

where the last relation holds when $|a_i|, q_i^2 \ll 1$.

Equation 2.11 describes the ion's trajectory as a slow harmonic oscillation at the sec-

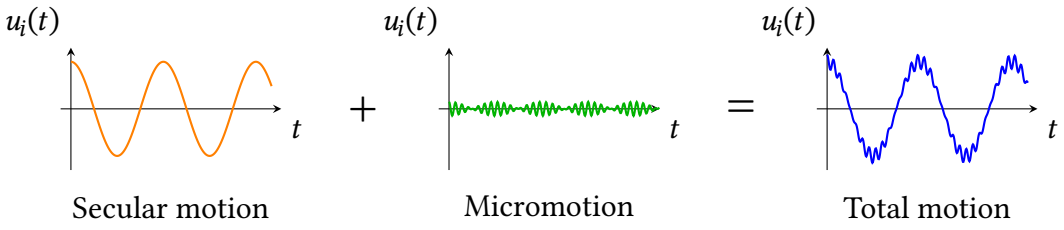

FIGURE 2.3: Visualization of ion motion in a linear Paul trap, with displacement along a generic direction denoted by u_i. The total motion (right) results from the superposition of a slow secular oscillation (left) at the secular frequency ω_i, and a fast micromotion (middle) at the RF drive frequency Ω_{RF}.

ular frequency ω_i, called *secular motion*, modulated by a fast component at the drive frequency Ω_{RF}, called *micromotion*, depicted in Figure 2.3.

This radial confinement is the fundamental mechanism of the trap; full three-dimensional confinement is achieved by combining this radial potential with axial DC confinement, discussed in the next section.

2.1.3 Axial confinement

Full three-dimensional confinement in a linear Paul trap is achieved by adding two DC endcap electrodes along the z-axis, as illustrated in panel (b) of Figure 2.1. This addition modifies the symmetry relations in Equation 2.6 to:

$$\begin{cases} \alpha_x = \alpha_y = -\alpha, & \alpha_z = 2\alpha, \\ \beta_x = -\beta_y = \beta, & \beta_z = 0, \end{cases} \tag{2.13}$$

which still satisfy Laplace's equation. The radial RF confinement remains essentially unchanged, although the radial stability parameters now satisfy $q_x = -q_y$ and the corresponding stability diagrams differ slightly from the purely radial case.

Along the axial direction, the ion experiences a simple harmonic potential due to the DC endcaps and its motion is governed by:

$$\left[m\frac{\mathrm{d}^2}{\mathrm{d}t^2} + 2Q\alpha\right] z = 0, \tag{2.14}$$

which corresponds to a one-dimensional harmonic oscillator with axial secular frequency:

$$\omega_z = \sqrt{\frac{2Q\alpha}{m}}. \tag{2.15}$$

In combination with the radial confinement previously introduced, this produces the full three-dimensional trapping potential that stabilizes the ion in all spatial directions. For storage of ion strings, the axial frequency is required to be smaller than the radial frequencies, $\omega_z < \omega_x, \omega_y$. Typical operating values for them are axial frequencies $\omega_z \sim 2\pi \times 1\,\mathrm{MHz}$ and radial frequencies $\omega_x, \omega_y \sim 2\pi \times 2\,\mathrm{MHz}$ to $5\,\mathrm{MHz}$ [37, 56].

2.1.4 Pseudopotential approximation

As discussed in the previous sections, in a linear Paul trap the ion experiences a rapidly oscillating RF field in the radial plane, superimposed on a static potential along the axial direction. The resulting motion in the radial plane can be decomposed into a slow *secular motion* at the radial secular frequency and a fast *micromotion* at the RF drive frequency Ω_{RF}. While this description captures the full time-dependent dynamics, it is often convenient to introduce an effective, *time-averaged* potential that reproduces the net confining effect of the RF field while neglecting the rapid micromotion oscillations. This approximation is valid when $\omega_i \ll \Omega_{\mathrm{RF}}$, so that the micromotion amplitude remains small compared to the secular motion. The resulting effective potential is commonly referred to as the *pseudopotential* or *ponderomotive potential* [52, 57, 58].

Formally, the time-average operation over one RF period, $T = 2\pi/\Omega_{\mathrm{RF}}$, is defined as:

$$\langle \vec{A} \rangle_T = \frac{1}{T} \int_0^T \vec{A}(t)\, dt\ , \tag{2.16}$$

which removes the rapidly oscillating micromotion terms and leaves only the slowly varying contributions responsible for the net confinement. For a generic RF electric field $\vec{E}_{\mathrm{RF}}(\vec{r}, t) = -\vec{\nabla} V_{\mathrm{RF}}(\vec{r}, t)$, the resulting pseudopotential is:

$$\Phi_{\mathrm{PS}}(\vec{r}) = \frac{Q^2}{2m\Omega_{\mathrm{RF}}^2} \left\langle \left| \vec{E}_{\mathrm{RF}}(\vec{r}, t) \right|^2 \right\rangle_T . \tag{2.17}$$

In a linear Paul trap, the RF field in the radial plane can be expressed as

$$\vec{E}_{\mathrm{RF}}(\vec{r}, t) = \vec{E}_0(\vec{r}) \cos(\Omega_{\mathrm{RF}} t), \tag{2.18}$$

where $\vec{E}_0(\vec{r})$ is the field amplitude at position $\vec{r} = (x, y, z)$. Thus, the pseudopotential becomes:

$$\Phi_{\mathrm{PS}}(\vec{r}) = \frac{Q^2}{4m\Omega_{\mathrm{RF}}^2} |\vec{E}_0(\vec{r})|^2, \tag{2.19}$$

effectively describing confinement in the radial plane. This represents an effective static potential, often referred to as *quasi-static* because it arises from averaging over the rapid RF oscillations, capturing only the slow secular motion. The full three-dimensional confinement of the ion is then obtained by combining this radial pseudopotential with the static axial potential provided by the DC endcaps Φ_{DC}, so that the total potential is:

$$\Phi\left(\vec{r}\right) = \Phi_{\mathrm{PS}}\left(\vec{r}\right) + \Phi_{\mathrm{DC}}\left(\vec{r}\right), \tag{2.20}$$

Consequently, the ion is localized near the minimum of the pseudopotential, which coincides with the *RF null* $(x, y = 0)$ along the trap axis.

The secular frequencies of the ion, corresponding to the slow harmonic motion, can be obtained from the curvature of the total potential:

$$\omega_{x,y} = \sqrt{\frac{Q}{m}\frac{\partial^2 \Phi_{\mathrm{PS}}}{\partial x^2}}, \qquad \omega_z = \sqrt{\frac{2Q\alpha}{m}}, \tag{2.21}$$

where α characterizes the static axial confinement.

It is important to emphasize that the pseudopotential is an approximation, valid only when the RF drive frequency Ω_{RF} is much larger than the secular frequencies ω_i. The approximation, however, breaks down if the ion is significantly displaced from the RF null, for example by stray electric fields, in which case the instantaneous RF field must be considered to accurately describe the ion dynamics [27, 28].

Beyond providing physical intuition, the pseudopotential approximation greatly simplifies numerical simulations. By replacing the full time-dependent RF field with a quasi-static one, it becomes possible to extract the main trapping quantities, such as the pseudopotential landscape and the secular frequencies, with simulations based only on electrostatics. This treatment forms the basis of the trap simulations performed throughout this thesis, enabling a quantitative characterization of the system while avoiding the computational complexity of solving the full time-dependent equations of motion.

2.2 Surface-electrode traps

In the previous section, we discussed the macroscopic linear Paul trap, which represents the fundamental architecture for ion confinement and manipulation. While these traps provide excellent optical access and long trapping lifetimes, they are not considered a scalable platform for large-scale quantum computation [27, 28]. The main limitation

arises from their 3D geometry: extending the design to include multiple trapping zones or junctions in a 3D configuration rapidly increases mechanical complexity, optical obstruction, and fabrication difficulty. As a result, implementing large-scale, interconnected ion networks in traditional 3D traps becomes difficult [34, 59].

The limitations of macroscopic linear Paul traps have motivated a fundamental shift in trap design, leading to the development of *surface-electrode traps* (or *planar traps*) [36, 37, 38, 60]. In contrast to conventional three-dimensional Paul traps, where electrodes surround the ions in space, in surface traps the RF and DC electrodes are deposited on a flat substrate, forming a two-dimensional geometry; the resulting electric fields create a pseudopotential minimum above the surface, where the ions are stably confined. This design marks a conceptual transition from bulky and mechanically assembled 3D devices to a fully *2D microfabricated architecture*. The planar configuration offers several important advantages over traditional macroscopic designs:

- **Fabrication:** planar traps can be produced using standard microfabrication processes, similar to those employed in semiconductor manufacturing, allowing precise electrode definition and reproducible fabrication.
- **Scalability:** the 2D geometry naturally supports complex, interconnected electrode layouts with multiple trapping zones, transport channels, and junctions integrated on a single chip.
- **Integration of electronic and optical components:** planar fabrication techniques enable the integration of control electronics, photodetectors, waveguides, and other optical elements directly on the chip, facilitating compact, modular, and highly controllable quantum devices.

These features make surface-electrode traps a promising and scalable platform for next-generation trapped-ion quantum processors.

2.2.1 Confinement in surface-electrode traps

The operation of a surface-electrode trap relies on the same fundamental principles as the linear Paul trap: a radio-frequency voltage applied to specific electrodes generates a rapidly oscillating electric field, whose net effect on the ion can be described using the *pseudopotential approximation*, while static voltages applied to other electrodes provide additional control over the axial confinement and allow fine-tuning of the ion's equilibrium position. Nevertheless, unlike the ideal quadrupole field produced in a macroscopic linear Paul trap, the electric field generated by planar electrodes in a surface trap

is inherently distorted by the open geometry and the finite extent of the electrodes. As a result, the resulting pseudopotential deviates from the perfect quadrupole form, and its curvature depends on the electrode layout [36, 37].

Figure 2.4 illustrates the typical *five-wire* surface trap geometry, one of the most widely adopted designs [38, 60]. The name refers to the five principal electrodes visible in the central region: two RF rails placed symmetrically about the trap axis and three DC electrodes (a central one flanked by two outer electrodes). Panel (a) shows how this planar arrangement can be viewed as the projection of a conventional 3D linear Paul trap onto a single plane, while panel (b) depicts the corresponding pseudopotential landscape above the surface. The minimum of this pseudopotential defines the *RF null*, located at a height h where the ion experiences no driven micromotion; for surface-electrode traps, the ion-surface distance is typically tens to hundreds of micrometers. As in the 3D case, the local curvature of the total potential around this minimum determines the secular frequencies ω_x, ω_y, and ω_z, which characterize the effective confinement along the radial and axial directions [37, 36]. In planar geometries, however, the asymmetry of the field distribution can lead to unequal radial frequencies and to a reduced overall confinement strength for a given RF voltage.

Another key parameter is the *trap depth*, defined as the potential energy barrier between the pseudopotential minimum and the nearest escape point, or *saddle point*, in the potential landscape [38]. In surface traps, this saddle point typically lies along a direction tilted toward the electrode plane, reflecting the open structure of the field. Consequently, the achievable trap depth is generally smaller than in macroscopic 3D

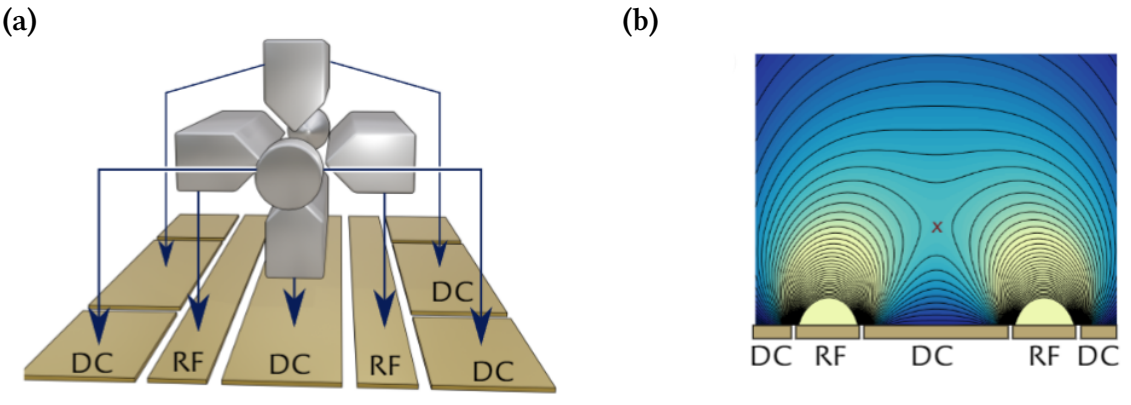

FIGURE 2.4: Schematic of a five-wire surface-electrode trap. Panel (a) shows the correspondence between a conventional 3D linear Paul trap and the planar 2D electrode layout. Panel (b) depicts the resulting pseudopotential above the surface, with the minimum (RF null) located at height h, where the ion is confined. Image courtesy of Martin van Mourik, adapted from his thesis [47].

traps and this can reduce trapping lifetimes.

Despite these geometric differences, the core physical mechanisms remain the same: the interplay between RF and DC fields creates a three-dimensional confining potential whose curvature sets the secular frequencies, while the pseudopotential minimum defines the ion equilibrium position. In the following section, the tools to simulate surface-electrode traps are presented.

2.2.2 Simulation methods for surface-electrode traps

The main goal of this thesis is to investigate ion heating caused by RF technical noise in surface-electrode traps. Achieving this requires simulating the pseudopotential landscape for arbitrary trap configurations, which is made possible through numerical simulations based on established analytical results from the literature. In the following, we introduce a model based on the gapless-plane approximation, which provides the theoretical foundation for the simulations presented in this work.

Gapless-plane approximation

The *gapless-plane approximation* idealizes the electrode surface as a continuous and perfectly conducting plane, where each electrode is held at a uniform potential, i.e. the potential is constant across the electrode, and no gaps exist between neighboring electrodes. This simplification captures the leading-order behavior of the electric field above the trap while neglecting finite gaps, surface roughness, and other fabrication imperfections [61, 62, 63]. In typical microfabricated surface traps, the electrode gaps are on the order of $1 - 10\,\mu\text{m}$, which is much smaller than the ion–surface distance. As a result, this approximation provides accurate estimates of the electric field and pseudopotential at the ion position, while greatly simplifying analytical calculations.

Within the gapless-plane approximation, the calculation of the electric field above an arbitrarily shaped planar electrode can be formulated as a Dirichlet problem in electrostatics. Consider a flat region in the x–y plane held at a fixed potential V, while the rest of the plane is grounded, and let $\mathscr{C}$ denote the closed boundary of the electrode (see Figure 2.5). The electric field at a point $\vec{x}$ above the plane is given by:

$$\vec{E}(\vec{x}) = -\vec{\nabla}\Phi(\vec{x}), \tag{2.22}$$

where $\Phi(\vec{x})$ is the electrostatic potential. Using Green's theorem and the method of images to satisfy the boundary conditions, the potential above the plane can be expressed

solely in terms of the electrode surface:

$$\Phi(\vec{x}) = \frac{V}{2\pi}\,\Omega(\vec{x}), \tag{2.23}$$

where $\Omega(\vec{x})$ is the solid angle subtended by the electrode as seen from $\vec{x}$, with a sign convention such that $\Omega > 0$ if the electrode normal points away from the observation point. By taking the gradient of the solid angle and applying a standard vector identity, the electric field can be rewritten as a line integral along the electrode boundary $\mathscr{C}$:

$$\vec{E}(\vec{x}) = \frac{V}{2\pi}\oint_{\mathscr{C}} \frac{(\vec{x}-\vec{x}')\times \mathrm{d}\vec{s}'}{|\vec{x}-\vec{x}'|^3}, \tag{2.24}$$

where $\mathrm{d}\vec{s}'$ is an infinitesimal segment along the contour, oriented according to the right-hand rule relative to the outward electrode normal. This expression is directly analogous to the Biot-Savart law in magnetostatics, with the cross product reversed to yield the correct electric field direction.

The main advantage of this formulation is that the field produced by a trap electrode in space can be computed using only its boundary contour, without integrating over the entire electrode surface. This greatly simplifies the numerical computation of the electric field for arbitrary planar electrode shapes [64] and provides the foundation for efficient numerical simulations of surface-electrode traps [63].

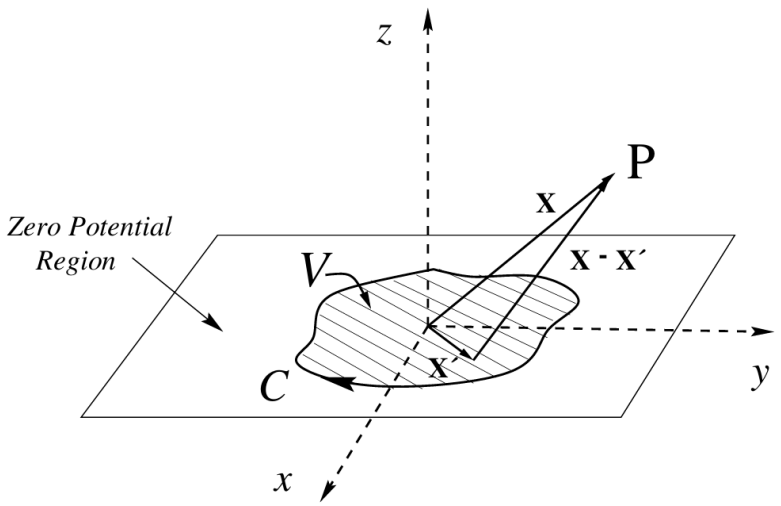

FIGURE 2.5: Schematic of the gapless-plane approximation. A charged flat plate is located in the x–y plane, bounded by the curve $\mathscr{C}$, and held at a fixed potential V. The potential is zero in the surrounding region of the plane ($z = 0$) outside $\mathscr{C}$. The source point on the electrode is denoted by $\vec{x}'$, while the observation point above the plane is $\vec{x}$, referred to as P [63].

Simulations in Python

To put the analytical framework described above into practice, all simulations presented in this work are performed using the ELECTRODE Python package, which implements the gapless-plane approximation to model arbitrary electrode geometries [65]. This framework provides a flexible and efficient tool for computing the resulting electric field and pseudopotential landscapes, as well as key quantities such as the secular frequencies and the trap depth, offering a complete characterization of the RF confinement.

In particular, the ELECTRODE package computes the potential generated by an arbitrary planar electrode by evaluating the line-integral expression of Equation 2.24. Each electrode boundary is discretized into a collection of short line segments, and the contribution of every segment to the electric field is evaluated analytically using the kernel derived from the gapless-plane formulation. This approach relies solely on the electrode contours, without meshing the electrode surfaces, and therefore preserves the numerical efficiency that motivates the gapless-plane approximation. Once the RF field $\vec{E}_{\mathrm{RF}}(\vec{r})$ is obtained on a spatial grid, the pseudopotential is computed using Equation 2.19, and relevant trap parameters such as trap depth, ion–surface distance, and secular frequencies are extracted from the curvature and local structure of the resulting potential landscape.

This approach effectively combines analytical and numerical techniques. It preserves the intuitive understanding provided by the analytical model while extending its applicability to arbitrary electrode geometries and voltage configurations. Because the calculations are quasi-static, the method remains computationally efficient, making it well suited for performing broad parameter scans and design optimizations. These capabilities form the basis for the simulation results presented in the following chapters.

2.2.3 Simulation of DC control potentials

A final remark concerns the *DC control electrodes*, which play a crucial role in the operation of surface-electrode ion traps. These DC electrodes are used to finely tune the static potential in the trapping region. Their purpose extends beyond providing axial confinement: they enable precise positioning of the ion within the RF null and allow for the controlled shaping of the overall trapping potential landscape.

In ideal conditions, an ion resides exactly at the RF null, where the pseudopotential reaches its minimum and the micromotion amplitude reaches its smallest possible value, given by $A_i q_i/2$ (see Equation 2.11). In real trapping conditions, however, stray electric fields can displace the ion from the RF null, resulting in a micromotion amplitude larger

than the ideal value. The resulting *excess micromotion* not only increases the ion's kinetic energy but also couples to the secular motion, leading to excess motional heating, reduced laser cooling efficiency, and frequency shifts in high-precision spectroscopic measurements [28, 66, 67]. To counteract these effects, DC voltages applied to the control electrodes generate compensating fields that restore the ion to the RF null, a procedure collectively known as *micromotion compensation*. These same electrodes also set the static confinement along the axial direction, allowing the axial secular frequency to be tuned to a desired value, typically around 1 MHz for standard operating conditions [27]. This dual role of the DC electrodes makes an accurate electrostatic model essential for relating applied voltages to both the local electric field and the curvature of the trapping potential at the ion position.

To establish this relation, we note that the voltages applied to the trap electrodes generate an electrostatic potential that can be expressed as a linear combination of orthogonal basis functions with well-defined spatial symmetries. A natural and widely adopted choice for this basis is the set of *spherical harmonic potentials* [62], which span all charge-free solutions of Laplace's equation up to a chosen order. This means that each harmonic $Y_{l,n}$, defined by its degree l and order n, satisfies Laplace's equation:

$$\nabla^2 Y_{l,n} = 0, \tag{2.25}$$

and an expansion up to $l = 2$ is sufficient to capture all quadratic contributions relevant to ion trapping and confinement [62].

The experimental realization of these multipole components, can be done computing a set of electrode voltages, referred to as the *shim set*, that reproduces each spherical harmonic potential. The computation proceeds as follows: in the simulation, each electrode k is sequentially set to 1 V, while all the others are grounded, and the resulting potential $\Phi_k(x_i, y_i, z_i)$ is sampled over a cubic grid centered on the chosen trapping location along the trap axis. The extracted potential is then fitted as a linear combination of n_m spherical harmonic functions:

$$\sum_m Y_m(x_i, y_i, z_i)\, w_m^{(k)} = \Phi_k(x_i, y_i, z_i), \tag{2.26}$$

where $w_m^{(k)}$ are the expansion coefficients. Expressing it in matrix form, the coefficients are obtained using a least-squares fit via the *pseudoinverse* of Y [68]:

$$Y\,\vec{w}^{(k)} = \vec{\Phi}_k \quad \rightarrow \quad \vec{w}^{(k)} = Y^{-1}\vec{\Phi}_k. \tag{2.27}$$

Each vector $\vec{w}^{(k)}$ quantifies how electrode k contributes to the different harmonic components. Stacking all $\vec{w}^{(k)}$ into the matrix W, the electrode voltages $\vec{V}_m$ that reproduce a given harmonic potential are obtained from:

$$WV = \mathbb{1}, \qquad V = W^{-1}. \tag{2.28}$$

Terms that do not contribute to the electric field (e.g. $Y_{0,0}$) can be omitted to minimize voltage amplitudes and simplify the solution space.

The resulting matrix V defines the *shim set*: a direct linear mapping between electrode voltages and spherical harmonic potentials. It can be formally shown that for a given desired ion position and set of motional frequencies, there exists a unique linear combination of spherical harmonic potentials that reproduces these conditions [47]. In practice, this means that by appropriately adjusting the voltages of the control electrodes, one can fully shape the trapping potential: setting the ion's equilibrium position, tuning the secular frequencies along each axis, and compensating stray fields. This capability also extends to dynamic operations such as shuttling ions along the trap or changing the orientation of ion chains, all achieved without altering the RF drive. Consequently, the shim set provides a direct and systematic mapping between electrode voltages and the experimentally relevant trapping parameters, offering complete control over the static potential landscape.

2.3 Qubits in trapped ions

The previous sections have provided a detailed overview of how ions can be confined and controlled using electromagnetic fields. Building on this foundation, we now explore how such trapped ions serve as fundamental units for quantum information processing.

In classical information theory, the fundamental unit of information is the *bit*, which can take one of two values, 0 or 1. Its quantum analogue is the *quantum bit*, or *qubit*, which requires a physical system with exactly two well-defined energy levels. These two levels, denoted $|0\rangle$ and $|1\rangle$, form the computational basis of the qubit. Unlike a classical bit, a qubit can exist in a coherent superposition of these states, allowing it to occupy both $|0\rangle$ and $|1\rangle$ simultaneously. Formally, the most general state of a single qubit undergoing unitary evolution can be written as

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle, \tag{2.29}$$

where α and β are complex coefficients constrained by normalization:

$$|\alpha|^2 + |\beta|^2 = 1. \tag{2.30}$$

Upon measurement in the computational basis, the qubit collapses to $|0\rangle$ with probability $|\alpha|^2$ or to $|1\rangle$ with probability $|\beta|^2$, which underlines the inherently probabilistic nature of the qubits.

Single trapped ions provide a natural realization of qubits. Two long-lived internal electronic states of the ion serve as the qubit levels $|0\rangle$ and $|1\rangle$, forming an almost ideal two-level system. These states can be coherently manipulated using laser fields, and their long coherence times and well-resolved optical transitions enable high-fidelity quantum operations. The coupling between the ion and laser light, which governs both coherent control and state readout, is extensively discussed in the literature [28, 67], and we summarize here the essential concepts.

2.3.1 The $^{40}Ca^{+}$ ion

A common choice for trapped-ion quantum computing is the $^{40}Ca^{+}$ ion, due to its convenient level structure and accessible optical transitions. In this system, the qubit is often encoded in the ground state $4S_{1/2}$ as $|0\rangle$ and the metastable excited state $3D_{5/2}$ as $|1\rangle$, which has a lifetime on the order of a second. Actually, the internal structure of the ion is more complex. The $4S_{1/2}$ level is split into two Zeeman sublevels, and the $3D_{5/2}$ level into six Zeeman sublevels. When a static magnetic field is applied, these states are energetically separated by the Zeeman effect, lifting their degeneracy. This splitting makes it possible to spectrally isolate and address a single, well-defined transition between a specific pair of Zeeman sublevels. By tuning the laser frequency and polarization accordingly, only this selected transition is driven, while all others remain far off-resonant. Under these conditions, the system can be accurately treated as an effective two-level system forming the physical qubit.

Auxiliary levels are employed for laser cooling, optical pumping, and state-dependent fluorescence detection. As shown in panel (a) of Figure 2.6, the relevant transitions are driven by lasers at well-defined wavelengths: the $4S_{1/2} \leftrightarrow 4P_{1/2}$ transition near 397 nm is used for Doppler cooling and fluorescence detection, the $3D_{3/2} \leftrightarrow 4P_{1/2}$ transition at 866 nm repumps population from the metastable $3D_{3/2}$ state, and the narrow $4S_{1/2} \leftrightarrow 3D_{5/2}$ electric quadrupole transition at 729 nm enables coherent qubit manipulation. Additional transitions can be engineered to implement more precise control.

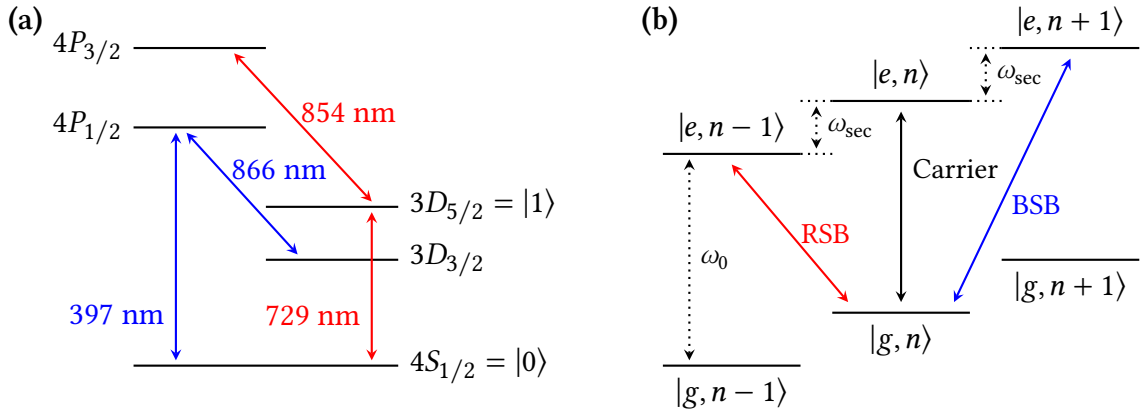

Figure 2.6: Energy levels and sideband transitions in $^{40}Ca^+$. (a) Simplified level structure of $^{40}Ca^+$ showing the qubit states $|0\rangle = 4S_{1/2}$ and $|1\rangle = 3D_{5/2}$. The transitions used for Doppler cooling and qubit manipulations are highlighted blue and red, respectively. (b) Carrier and first-order sideband transitions derived from the Lamb–Dicke approximation.

2.3.2 Ion-light interaction

As previously introduced, an ion interacting with a narrow-band coherent laser field can, in most situations, be accurately approximated as a two-level system. The atomic Hamiltonian is then:

$$\hat{H}_A = \frac{\hbar\omega_0}{2}\left(|e\rangle\langle e| - |g\rangle\langle g|\right) = \frac{\hbar\omega_0}{2}\,\sigma_z, \tag{2.31}$$

where ω_0 is the energy splitting between the electronic ground state $|g\rangle$ and excited state $|e\rangle$ [28]. The mapping to the Pauli matrices is given by:

$$\sigma_x = |g\rangle\langle e| + |e\rangle\langle g|, \quad \sigma_y = i\left(|g\rangle\langle e| - |e\rangle\langle g|\right), \quad \sigma_z = |e\rangle\langle e| - |g\rangle\langle g|. \tag{2.32}$$

We assume that the ion is trapped in a harmonic potential, valid under the pseudopotential approximation (see Subsection 2.1.4). Considering one spatial dimension, the motional degree of freedom is described by a quantum harmonic oscillator:

$$\hat{H}_M = \hbar\omega_{\text{sec}}\left(\hat{n} + \frac{1}{2}\right), \tag{2.33}$$

where ω_{sec} is the secular frequency, $\hat{n} = \hat{a}^\dagger\hat{a}$ is the number operator, and $\hat{a}^\dagger$, $\hat{a}$ are the creation and annihilation operators. The total unperturbed Hamiltonian is then:

$$\hat{H}_0 = \hat{H}_A + \hat{H}_M\,. \tag{2.34}$$

The interaction between a laser field with wave number k, frequency ω_L and phase

θ_L, and the ion can be written in a general form with Hamiltonian:

$$\hat{H}_{\mathrm{LA}} = \frac{\hbar\Omega}{2}(\sigma_+ + \sigma_-)\left[e^{i(kx-\omega_L t+\theta_L)} + \text{c.c.}\right], \tag{2.35}$$

where

$$\sigma_+ = \frac{1}{2}(\sigma_x + i\sigma_y) = |e\rangle\langle g|, \qquad \sigma_- = \frac{1}{2}(\sigma_x - i\sigma_y) = |g\rangle\langle e|, \tag{2.36}$$

are the coupling operators between ground and excited states in the spin-1/2 algebra. The Rabi frequency Ω determines the laser coupling strength and is proportional to the transition matrix element $\langle g|\hat{H}_{\mathrm{LA}}|e\rangle$, which accounts for the specific Zeeman sublevels involved in the transition, as well as the electric-field amplitude E_0.

The laser-induced dynamics are conveniently described in the interaction picture, where the position operator $\hat{x}$ is expressed in terms of the motional operators via the Lamb–Dicke parameter η:

$$\hat{x} = \eta(\hat{a} + \hat{a}^\dagger), \qquad \eta = k\sqrt{\frac{\hbar}{2m\omega_{\mathrm{sec}}}}\cos\alpha, \tag{2.37}$$

with m being the ion mass and α the angle between the laser wave vector and the motional axis. Applying this and using the *rotating-wave approximation* (RWA), which neglects fast-oscillating off-resonant terms [27, 28], gives:

$$\hat{H}_{\mathrm{int}} \approx \frac{\hbar\Omega}{2}\,\sigma_+ \exp\left[i\eta(\hat{a}e^{-i\omega_{\mathrm{sec}}t} + \hat{a}^\dagger e^{i\omega_{\mathrm{sec}}t})\right]e^{-i(\delta t-\theta_L)} + \text{h.c.}, \tag{2.38}$$

where $\delta = \omega_L - \omega_0$ is the laser detuning from the atomic transition.

In most experimental conditions, the ion is confined within a region that is much smaller than the laser wavelength. We refer to this condition as the *Lamb–Dicke regime*, characterized by $\eta^2(2\bar{n}+1) \ll 1$, where $\bar{n} = \langle\hat{n}\rangle$ is the mean motional occupation number. In this regime, the exponential in Equation 2.38 can be expanded to first order:

$$\hat{H}_{\mathrm{int}} \approx \frac{\hbar\Omega}{2}\,\sigma_+\left[1 + i\eta(\hat{a}e^{-i\omega_{\mathrm{sec}}t} + \hat{a}^\dagger e^{i\omega_{\mathrm{sec}}t})\right]e^{-i(\delta t-\theta_L)} + \text{h.c.}. \tag{2.39}$$

It is important to notice that Equation 2.39 contains three resonant terms, corresponding to the transitions illustrated in panel (b) of Figure 2.6:

1. **Carrier** ($\delta = 0$): Couples internal states without changing the motional quantum number, $|g,n\rangle \leftrightarrow |e,n\rangle$, with Rabi frequency Ω:

$$\hat{H}_{\mathrm{car}} = \frac{\hbar\Omega}{2}\left(\sigma_+ e^{i\theta_L} + \sigma_- e^{-i\theta_L}\right). \tag{2.40}$$

2. **Red sideband** ($\delta = -\omega_{\text{sec}}$): Couples $|g,n\rangle \leftrightarrow |e,n-1\rangle$, removing one motional quantum, with strength $\Omega_{n,n-1} = \eta\Omega\sqrt{n}$:

$$\hat{H}_{\text{rsb}} = \frac{\hbar\eta\Omega}{2}\left(\hat{a}\sigma_+ e^{i\theta_L} + \hat{a}^\dagger\sigma_- e^{-i\theta_L}\right). \tag{2.41}$$

3. **Blue sideband** ($\delta = +\omega_{\text{sec}}$): Couples $|g,n\rangle \leftrightarrow |e,n+1\rangle$, creating one motional quantum, with strength $\Omega_{n,n+1} = \eta\Omega\sqrt{n+1}$:

$$\hat{H}_{\text{bsb}} = \frac{\hbar\eta\Omega}{2}\left(\hat{a}^\dagger\sigma_+ e^{i\theta_L} + \hat{a}\sigma_- e^{-i\theta_L}\right). \tag{2.42}$$

When the linewidth γ of the optical transition and the Rabi frequency Ω both satisfy $\gamma, \Omega \ll \omega_{\text{sec}}$, the motional sidebands are spectrally resolved, which enables individual addressing of the carrier and sideband transitions. This offers a way to manipulate the electronic state of the ion conditioned on its motional state and viceversa, which is at the base of the implementation of multi-qubit quantum gates in trapped ion quantum information processors.

2.4 Motional heating

The concepts introduced throughout this chapter, the principles of ion confinement in Paul traps, the advantages of surface-electrode geometries, and the interaction between trapped ions and laser fields, serve to establish the framework necessary to understand the main topic of this thesis, which constitutes also one of the most critical limitations in modern trapped-ion quantum technologies: *motional heating*.

In trapped-ion quantum computing, the quantized motion of the ions is not merely a secondary degree of freedom but an essential quantum resource. Collective motional modes act as a quantum bus that mediates interactions between individual qubits, enabling multi-qubit entangling gates such as the Cirac–Zoller [31] and Mølmer–Sørensen [28, 30] gates. For these operations to achieve high fidelity, the ions' motion must remain coherent; any uncontrolled coupling to the environment reduces gate fidelity and limits the scalability of trapped-ion quantum computers [27].

A central figure of merit in this context is the *motional decoherence*, which quantifies the loss of phase coherence of the ion's quantized motion due to fluctuating electric fields or other environmental perturbations. In the absence of decoherence, the motional state evolves coherently, preserving a well-defined phase relationship crucial for quantum logic operations. When motional decoherence occurs, this phase coherence is

disrupted, introducing random phase noise that degrades entangling gate fidelity and compromises reproducibility.

Motional heating is closely related to the decoherence of the ion motion: if the ion undergoes an unwanted change of its motional state, the coherence of its motion is disrupted. In other words, when the ion motion is 'heated" by environmental fluctuations, the system can no longer maintain coherent dynamics over long timescales. This effect is especially pronounced in surface-electrode traps, where ions are typically confined only tens of micrometres above the electrodes and are therefore more sensitive to electric-field noise coming from the trap surface, which is a direct responsible of the motional heating of the ion.

Understanding and mitigating motional heating is therefore a central goal in the design and optimization of surface-electrode traps. Since the heating rate directly determines the timescale of motional decoherence, it imposes a fundamental limit on how long an ion can store quantum information in its motion and on the fidelity of multi-qubit gate operations. The focus of this thesis is the study of motional heating induced by RF technical noise in surface-electrode traps. The theoretical modelling of such heating will be the topic of the following chapter.

3

Ion heating due to RF technical noise

In this chapter, we examine the ion motional heating resulting from *technical noise* on the RF drive in surface traps. The first section describes how electric field noise couples to the ion's motion, leading to heating, when resonant with the trap frequencies. The second section, instead, presents the theoretical derivation of ion heating models in presence of amplitude and phase noise on the RF drive.

3.1 Heating mechanism

As previously introduced, *motional heating* in ion traps refers to the unwanted increase of an ion's motional energy over time, caused by electric field noise from the environment [40, 42, 44]. It is described by the increase of the average phonon number of a specific motional mode over time, given by:

$$\bar{n} = \langle \hat{a}^\dagger \hat{a} \rangle, \tag{3.1}$$

where $\hat{a}$ and $\hat{a}^\dagger$ denote the annihilation and creation operators, respectively, for the chosen motional mode (see Equation 2.33). We refer to the rate of change of motional excitation over time as the *ion heating rate*, $\dot{\bar{n}}$. This quantity, indeed, captures how quickly the ion absorbs motional quanta due to external noise sources [44].

Different noise sources can contribute to the motional heating of trapped ions. When considering planar traps, we can differentiate between two main categories:

- *Surface noise (anomalous heating):* Electric field noise originating from 'microscopic' surface effects on the trap electrodes, such as fluctuating patch potentials or dipoles (small with respect to the dimension of the trap electrodes). Its strength typically increases sharply as the ion–surface distance decreases, scaling approximately as d^{-4} [40].

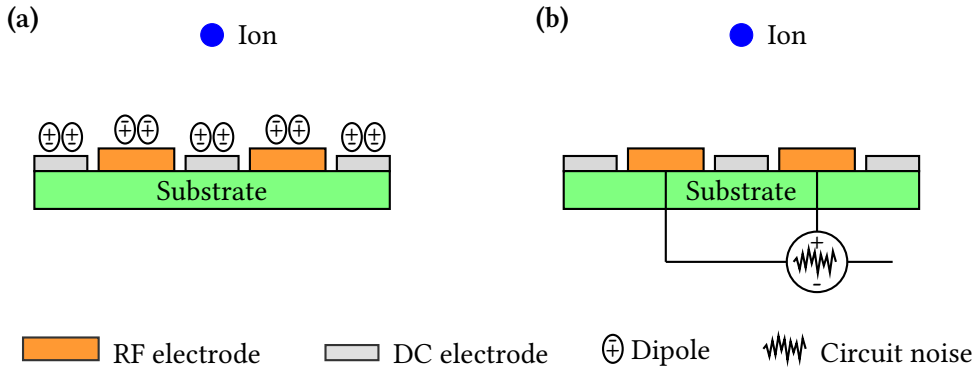

Figure 3.1: Schematic comparison of two motional heating mechanisms in surface-trap ions, showing the electric field at the trap minimum: (a) Surface-induced noise from fluctuating dipoles on the electrode surface, superimposed on the static RF field, which results in an oscillating electric field. (b) Technical voltage noise from the RF drive circuitry, which directly causes oscillations of the electric field.

- *Technical noise:* Electric field noise arising from imperfections in the electronic circuitry driving the trap. It can appear on both the DC signals and the RF signals, and is directly injected through the electrodes, causing heating when resonant with the ion's motion.

Figure 3.1 illustrates the conceptual difference between the two distinct heating mechanisms. Panel (a) describes how surface noise arises from microscopic fluctuating dipoles on the trap electrodes. The field produced by these fluctuating patch potentials superimposes to the trapping field and can excite the ion motion. In panel (b), it is shown that technical noise originates from imperfections in the trap circuitry. In this case, the electric field at the ion position oscillates due to voltage fluctuations on the electrodes, which couple to the ion's motion and result in heating.

As both surface-induced heating [40, 42, 43, 44] and heating induced by noise on the DC electrodes [27, 69] have been treated extensively in the literature, in this work we focus solely on ion heating arising from technical noise on the RF drive signals, which we will refer to as *RF technical noise* hereafter. In the following sections, we will first focus on the noise in electronics in general and then we'll move to the mathematical description of RF technical noise in surface-electrode ion traps.

3.1.1 Noise in electronics

Before discussing RF technical noise in ion traps, it is useful to introduce the general concept of electronic noise and some related quantities.

In electronic systems, the term *noise* denotes random fluctuations of physical quantities (such as voltage, current, or phase) that are superimposed on the desired signal. These fluctuations stem from the stochastic nature of microscopic processes; thus, their exact temporal evolution cannot be predicted. For this reason, noise is modeled as a *random process* $X(t)$, i.e. a family of random variables indexed by time. Consequently, to describe and analyze such signals, one must rely on statistical quantities rather than explicit time-domain trajectories [70].

A natural and widely used statistical tool for characterizing a random signal is the *autocorrelation function*, which measures the similarity of the process with a time-shifted version of itself. In other words, it quantifies how predictable the signal is over a given time delay τ. The autocorrelation is a real and even function defined as:

$$R_{XX}(\tau) = \int_{-\infty}^{+\infty} X(t)\, X(t+\tau)\, dt = \langle X(t)\, X(t+\tau) \rangle, \tag{3.2}$$

where $\langle \cdot \rangle$ denotes the ensemble average over all realizations of the process. This notation is commonly used because it provides a more compact and readable way to represent averages, avoiding repeated integral expressions. For a stationary random process, i.e. a process whose statistical properties do not change over time, $R_{XX}(\tau)$ depends only on the time difference τ and not on the absolute reference time t. At zero delay, the autocorrelation reduces to the mean-square value of the process:

$$R_{XX}(0) = \int_{-\infty}^{+\infty} X^2(t)\, dt = \left\langle X^2(t) \right\rangle = E, \tag{3.3}$$

where E is called the *energy* of the signal. It is clear that:

$$|R_{XX}(\tau)| \leq R_{XX}(0), \tag{3.4}$$

since the maximum similarity is reached when the two signals are identical ($\tau = 0$). If the process has zero mean, this energy coincides with the variance, and, in physical terms, corresponds to the average power carried by the fluctuations.

While the autocorrelation function characterizes the noise in the time domain, it is also useful to analyze how these random fluctuations are distributed across different frequencies. This leads to the concept of the *power spectral density* (PSD), which quantifies how the variance of the process is spread over frequency components. Formally, for a stationary random process, the Wiener–Khinchin theorem relates the PSD to the

autocorrelation function via a Fourier transform [71, 72]:

$$S_X(\omega) = \int_{-\infty}^{+\infty} R_{XX}(\tau)\, e^{-i\omega\tau}\, d\tau. \tag{3.5}$$

Conversely, the autocorrelation can be obtained from the inverse Fourier transform of $S_X(\omega)$, given by:

$$R_{XX}(\tau) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} S_X(\omega)\, e^{i\omega\tau}\, d\omega. \tag{3.6}$$

These relations demonstrate that the time-domain and frequency-domain descriptions of a random process are mathematically equivalent. In practice, the PSD is often more convenient than the time-domain description, as it directly reveals the spectral content of the noise. In systems containing resonant circuits, for example, only specific frequency components are significantly transmitted or amplified, while others are strongly attenuated. Thus, the PSD allows one to focus on and quantify only the frequencies that are relevant for the system's behavior. The connection between the PSD and the total noise power can be made explicit by considering the mean-square value of the process. For a stationary random process, the total variance can be expressed as the integral of the PSD over all frequencies:

$$\left\langle X^2(t) \right\rangle = R_{XX}(0) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} S_X(\omega)\, d\omega. \tag{3.7}$$

This relation shows that the PSD provides a direct measure of the contribution of each frequency band to the total noise power.

A particularly important theoretical model is that of *white noise*. Statistically, it is defined as a noise signal whose autocorrelation is proportional to a Dirac delta function:

$$R_{XX}(\tau) = \lambda\, \delta(\tau), \tag{3.8}$$

meaning that the autocorrelation vanishes for any nonzero time delay. In other words, the process is completely uncorrelated with itself at different times, making it entirely unpredictable. For the same reason, white noise has a flat spectrum, which arises directly from the absence of temporal correlation: with no predictable structure, every frequency component contributes equally to the total power. We can prove this simply computing the PSD of white noise, applying Equation 3.5:

$$S_N(\omega) = \int_{-\infty}^{+\infty} \lambda\, \delta(\tau)\, e^{-i\omega\tau}\, d\tau = \lambda, \tag{3.9}$$

which shows that the PSD of an ideal white-noise process is constant over all frequencies. Although such an ideal process cannot exist physically, since it would require infinite power due to its infinite bandwidth, this model provides a useful approximation for real noise signals within a limited frequency range. In systems containing resonant circuits, such as the RF drive in a Paul trap, the relevant bandwidth is typically very narrow compared to the full spectrum of the noise, so the noise within this band can be treated, without loss of generality, as white noise, with an approximately constant spectral density over the frequencies of interest.

This perspective naturally leads to a theoretical description of electronic noise based on two key concepts: the autocorrelation function, which characterizes temporal correlations, and the power spectral density, which describes how fluctuations are distributed across frequencies. In particular, the white-noise model, with its delta-correlated autocorrelation and flat spectral density, provides a convenient and accurate approximation for noise in narrowband systems. In the next section, these concepts will be applied to describe technical noise in ion traps.

3.1.2 Theoretical description

To describe the impact of RF technical noise on trapped ions, it is necessary to characterize how fluctuating forces influence the ion's motion. When an ion experiences a force resonant with one of its secular frequencies, $\omega_{x,y,z}$, the energy of the corresponding motional mode increases, leading to heating. The rate at which the mean phonon number grows, i.e. the heating rate, is given by the well-known relation [40]:

$$\dot{\bar{n}}_{x,y,z} = \frac{S_{F_N}(\omega_N)}{4m\hbar\omega_{x,y,z}}\,, \tag{3.10}$$

where $S_{F_N}(\omega_N)$ is the spectral density of the noise force at the noise frequency ω_N, i.e. the frequency component of the noise that drives the ion's motion. Depending on the noise mechanism, different noise components can excite the ion's motion along a specific direction, provided they are resonant with the secular frequency of the corresponding mode.

The force spectral density is directly connected to the spectral density of voltage fluctuations applied to the trap electrodes via:

$$\frac{S_{V_N}}{V_0^2} = \frac{S_{F_N}}{F_0^2}, \tag{3.11}$$

where V_0 is the amplitude of the RF trapping potential and $F_0^2 \equiv \left|\vec{F}_0\right|^2$ characterizes the strength of the electric force generated by the RF drive [41]. In the case of RF phase noise, the fluctuations are typically expressed as a phase noise spectral density $S_{\theta_N}(\omega_N)$ rather than a normalized voltage spectral density. This phase noise gives rise to an effective force spectral density at the ion position according to:

$$S_{F_N}(\omega_N) = S_{\theta_N}(\omega_N)\, F_0^2\,, \tag{3.12}$$

so that the corresponding heating rate can be evaluated, using Equation 3.10, without any additional normalization of the phase noise spectrum. Depending on the type of noise present in the experiment, one then uses either the voltage or phase noise spectral density to quantify the force fluctuations. In both cases, the spectral densities are computed according to the general definition given in Equation 3.5.

The net noise force originates from electric-field fluctuations at the ion's position. Using Newton's second law, it can be expressed in terms of the pseudopotential Φ_{PS} as

$$\vec{F}_N = -\vec{\nabla}U \approx -\vec{\nabla}\Phi_{\mathrm{PS}}\big|_{\vec{r}=\vec{r}_0} \propto -\vec{\nabla}|E|^2, \tag{3.13}$$

where $\vec{r}_0$ denotes the equilibrium position of the ion, and the last equality holds under the pseudopotential approximation (see Subsection 2.1.4. The magnitude of $\vec{F}_N$, and consequently the resulting heating rate, depends on the trap geometry and the spectral properties of the noise. This stochastic noise force should not be confused with a coherent, controlled drive applied to the ion, such as a laser used to implement quantum gates. While $\vec{F}_N$ leads to motional heating with a mean phonon number that grows linearly in time [40, 41], a resonant coherent force produces a quadratic increase in excitation [73]. The resulting difference in temporal behavior allows coherent excitation to be distinguished from stochastic noise, enabling the ion's motional heating arising from RF technical noise to be directly observed and quantified.

While Equation 3.13 provides the general framework linking electric-field fluctuations to a net noise force on the ion, the question remains as to which frequency components of the technical noise actually contribute to driving the ion's motion. Clarifying this point is essential before developing a quantitative description of technical noise in ion traps. Two main arguments must be kept in mind:

- Since the RF drive is delivered through a resonant circuit, broadband fluctuations outside the narrow resonance bandwidth are strongly suppressed, and only residual noise close to the carrier frequency reaches the trap electrodes.

- Within the pseudopotential approximation, efficient excitation of the ion motion occurs only when electric-field fluctuations are resonant with one of the secular frequencies $\omega_{x,y,z}$. Off-resonant components average out and do not lead to heating, as previously remarked.

Together, these considerations imply that the relevant contributions arise from the beat between the RF drive Ω_{RF} and noise components at frequencies $\Omega_{\mathrm{RF}} \pm \omega$, being ω a target motional mode. Thus, low frequency noise at frequency ω does not directly drive the motion; instead, mixing with the RF tone shifts it to the sidebands that can resonantly couple to the ion.

As a final remark, although the precise spectral characteristics of the electronic noise in the RF circuitry are unknown, the resonator bandwidth is narrow compared to Ω_{RF}. For this reason, the noise driving the resonator can be well approximated as white noise within the resonator bandwidth.

In summary, we have seen that ion heating is determined by the fluctuating force acting on the ion, which in turn is set by the local electric field at the ion's position. Identifying the relevant noise frequencies and their coupling to the secular motion therefore provides the essential link between technical noise at the trap electrodes and the resulting heating rates. To develop a quantitative description, we now turn to explicit noise models: starting from the analytic form of the electric field including noise contributions, we derive the corresponding pseudopotential, evaluate the resulting force on the ion, and obtain the associated heating rate.

3.2 Noise models

Before discussing specific models, it is useful to distinguish between two common types of technical noise affecting the RF drive: *amplitude noise* and *phase noise*.

- Amplitude noise refers to random fluctuations in the magnitude of the RF signal. These variations change the signal's envelope while its oscillation frequency remains unchanged.

- Phase noise refers to random fluctuations in the instantaneous phase of the RF signal. The amplitude stays constant, but the timing of the oscillations jitters around the ideal periodic evolution.

Although both types of technical noise can contribute to ion heating, their effects on the trap drive are qualitatively different. In particular, they are described by distinct

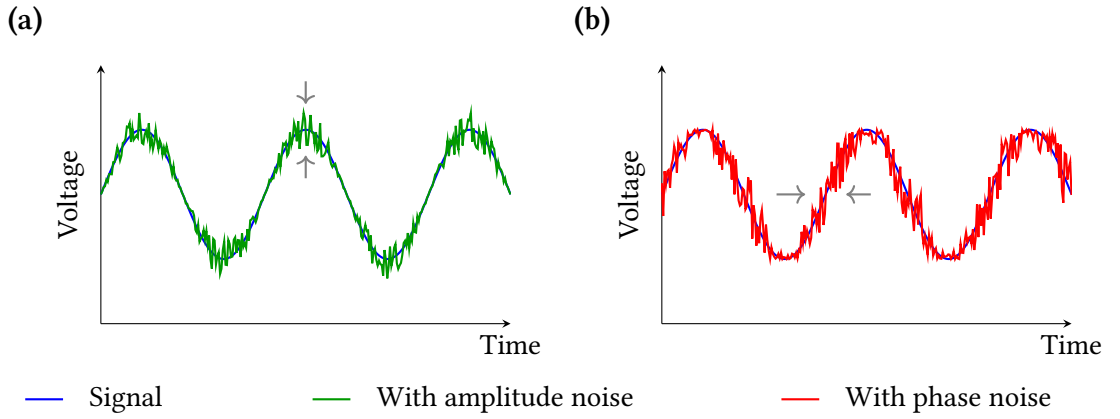

Figure 3.2: Conceptual comparison of amplitude and phase noise. (a) Amplitude noise: the sinusoidal signal experiences random modulation of its amplitude, most evident near the peaks. (b) Phase noise: the signal maintains constant amplitude, but its phase fluctuates randomly, leading to variations in the spacing of the oscillations.

mathematical models, which lead to different dependencies when calculating the corresponding heating rates. Figure 3.2 illustrates these differences in the time domain, when considering a random modulation on the signal. The gray arrows highlight the nature of the fluctuations: for amplitude noise they point along the vertical axis, indicating variations in the signal's magnitude, while for phase noise they point along the horizontal axis, indicating random shifts in the oscillation timing. These qualitative differences call for distinct mathematical descriptions of amplitude and phase noise.

In the following sections, we derive the corresponding models in detail, highlighting how the two types of noise lead to distinct heating rate models.

3.2.1 Amplitude noise model

We begin the analysis of the effect of amplitude noise considering a single RF drive signal. This simple case will later serve as the basis for the generalization to a trap driven by multiple RF signals.

As previously explained, we can express the drive voltage as the sum of the ideal RF signal and a noise contribution at frequency $\omega_N = \Omega_{\mathrm{RF}} \pm \omega$, where Ω_{RF} is the RF drive frequency and $|\omega| \ll \Omega_{\mathrm{RF}}$ [41]. For simplicity, and without loss of generality, we restrict the discussion to the component at $\Omega_{\mathrm{RF}} + \omega$. The resulting electric field, calculated at the ion position (x, y, z), takes the form:

$$\vec{E}_{\mathrm{RF}}(x, y, z, t) = \vec{E}_0(x, y, z)\left[\cos(\Omega_{\mathrm{RF}} t) + \chi_N \cos((\Omega_{\mathrm{RF}} + \omega)t)\right], \tag{3.14}$$

where $\chi_N \ll 1$ is the noise amplitude relative to the ideal field $\vec{E}_0(x,y,z)$. In the absence of noise ($\chi_N = 0$), Equation 2.17 leads to the standard pseudopotential expression:

$$\Phi_{\mathrm{PS}} = \frac{Q^2}{2m\Omega_{\mathrm{RF}}^2}\left\langle\left|\vec{E}_{\mathrm{RF}}(x,y,z,t)\right|^2\right\rangle_T = \frac{Q^2}{4m\Omega_{\mathrm{RF}}^2}\left|\vec{E}_0(x,y,z)\right|^2 , \tag{3.15}$$

where Q and m are the ion's charge and mass, and the time average is taken over one RF cycle, with period $T = 2\pi/\Omega_{\mathrm{RF}}$.

The additional noise term modulates the trapping field, producing a beat with the carrier that drives the ion at frequency ω, as discussed at the end of Subsection 3.1.2. To evaluate the resulting force via Equation 3.13, we first compute the pseudopotential at the ion position, starting from the squared magnitude of the electric field:

$$\begin{aligned}\left|\vec{E}_{\mathrm{RF}}\right|^2 = &\left|\vec{E}_0\right|^2\left[\cos^2(\Omega_{\mathrm{RF}}t) + \chi_N^2\cos^2((\Omega_{\mathrm{RF}}+\omega)t)\right] + \\ &+ \left|\vec{E}_0\right|^2\left[2\chi_N\cos((\Omega_{\mathrm{RF}}+\omega)t)\cos(\Omega_{\mathrm{RF}}t)\right]\end{aligned} \tag{3.16}$$

The first contribution represents the standard trapping potential; the second term is quadratic in the noise amplitude χ_N, while the third contribution is linear in χ_N. Since the noise amplitude $\chi_N \ll 1$, the quadratic term can be neglected, while the linear term is the relevant one, provided it survives after performing the time average operation over one RF cycle.

Before this operation, it is useful to apply Werner's identity to expand the product of cosines, $\cos(A)\cos(B) = \frac{1}{2}[\cos(A-B) + \cos(A+B)]$. This expansion produces in the cross-term a slow component oscillating at frequency ω and a fast component oscillating at frequency $2\Omega_{\mathrm{RF}} + \omega$. Since $|\omega| \ll \Omega_{\mathrm{RF}}$, all fast oscillating components oscillate at $\mathcal{O}(\Omega_{\mathrm{RF}})$ and vanish under the RF-cycle average, while the slow components vary negligibly over one RF cycle and can be taken outside the integral. Thus, after applying the time average operation, we get:

$$\left\langle\left|\vec{E}_{\mathrm{RF}}\right|^2\right\rangle = \left|\vec{E}_0\right|^2\left[\frac{1}{2} + \frac{\chi_N^2}{2} + \chi_N\cos(\omega t)\right] \approx \left|\vec{E}_0\right|^2\left[\frac{1}{2} + \chi_N\cos(\omega t)\right], \tag{3.17}$$

where in the last step we can neglect the quadratic term since $\chi_N \ll 1$. The resulting pseudopotential reads as:

$$\Phi_{\mathrm{PS}} = \frac{Q^2}{4m\Omega_{\mathrm{RF}}^2}\left\langle\left|\vec{E}_{\mathrm{RF}}\right|^2\right\rangle = \frac{Q^2\left|\vec{E}_0\right|^2}{4m\Omega_{\mathrm{RF}}^2}\left[\frac{1}{2} + \chi_N\cos(\omega t)\right] = \Phi_0 + \Phi_N . \tag{3.18}$$

Here, Φ_0 corresponds to the noiseless pseudopotential, while Φ_N represents the noise-induced modulation. Using Equation 3.13, we obtain the net force acting on the ion:

$$\vec{F}_N = -\frac{Q^2}{4m\Omega_{\mathrm{RF}}^2}\left[\vec{\nabla}\left(\left|\vec{E}_0\right|^2\right)\right]\chi_N\cos(\omega t) = \vec{F}_0\chi_N\cos(\omega t). \tag{3.19}$$

This is the oscillating force responsible for motional heating when resonant with the ion motion. The ion, however, resides at the RF null, defined as the point where the pseudopotential vanishes, as described in Section 2.2, and, consequently, $\vec{E} = 0$. Even though this, in general, does not imply $\nabla\vec{E} = 0$, in this specific case we can use the relation $\vec{\nabla}\left|\vec{A}\right|^2 = 2\left|\vec{A}\right|\left(\vec{\nabla}\left|\vec{A}\right|\right)$, thus $\vec{F}_N = 0$ if $\vec{E}_0 = 0$. Thus, that we can have a non-vanishing heating rate only if the electric field is not null at the ion position. We will see in Chapter 4 that this is the case for the X-junction in the QCCD architecture.

Combining Equation 3.10 with Equation 3.11, we can find an expression for ion heating in the presence of amplitude noise for a single RF drive:

$$\dot{\bar{n}} = \frac{F_0^2}{V_0^2}\frac{S_{V_N}(\omega_N)}{4m\hbar\omega} = \frac{Q^4}{16m^3\Omega_{\mathrm{RF}}^4\hbar\omega}\left[\vec{\nabla}\left(\left|\vec{E}_0\right|^2\right)\right]^2\frac{S_{V_N}(\Omega_{\mathrm{RF}}\pm\omega)}{V_0^2} \tag{3.20}$$

We can easily show that the spectral density can be evaluated at both the upper and lower RF sidebands, $S_{V_N}(\Omega_{\mathrm{RF}}\pm\omega)$, which are the frequency components that remain after the averaging operation and contribute to motional heating. Explicitly, we first compute the autocorrelation function, which, thanks to Equation 3.2, results in:

$$R_{V_N}(\tau) = V_0^2\left\langle\chi_N^2\right\rangle_T\left\langle\cos(\omega t)\cos(\omega(t+\tau))\right\rangle_T = \frac{V_0^2}{2}\left\langle\chi_N^2\right\rangle_T\cos(\omega\tau), \tag{3.21}$$

where only the constant term of the expansion of the product of cosine remains after the time average. Then, using Equation 3.5 with a generic noise frequency $\Omega_{\mathrm{RF}}\pm\omega_N$, the voltage spectral density is:

$$\begin{aligned} S_{V_N}(\Omega_{\mathrm{RF}}\pm\omega_N) &= V_0^2\left\langle\chi_N^2\right\rangle\int_{-\infty}^{+\infty}\left\langle\cos(\omega t)\cos(\omega(t+\tau))\right\rangle_T\ e^{-i(\Omega_{\mathrm{RF}}\pm\omega_N)\tau}d\tau \\ &= \frac{V_0^2}{2}\left\langle\chi_N^2\right\rangle\int_{-\infty}^{+\infty}\cos(\omega\tau)\ e^{-i(\Omega_{\mathrm{RF}}\pm\omega_N)\tau}d\tau \\ &= \frac{\pi V_0^2}{2}\left\langle\chi_N^2\right\rangle\left[\delta(\omega-\omega_N)+\delta(\omega+\omega_N)\right], \end{aligned} \tag{3.22}$$

which is not vanishing only for $\omega_N = \pm\omega$, as expected.

Since the noise spectrum is symmetric around the carrier, both sidebands give the same contribution, and in the following it is sufficient to consider only one of them to avoid redundancy. Moreover, the modulation amplitude χ_N explicitly enters the spectral density expression, so that stronger noise fluctuations directly enhance the heating rate.

The final equation for the heating rate, therefore, depends on two main factors, which determine how efficiently voltage fluctuations are converted into motional excitation of the trapped ion:

1. **Electrode geometry:** The gradient of the RF field, $\vec{\nabla}\left|\vec{E}_0\right|^2$, encodes the effect of the trap geometry and the spatial configuration of the electrodes. Steeper gradients lead to stronger coupling between voltage fluctuations and the ion's motion.
2. **Noise amplitude:** An increase in the spectral density S_{V_N} corresponds directly to an increase in the heating rate.

The single-drive analysis illustrates how amplitude fluctuations translate into a noise force and, ultimately, into motional heating. However, Paul traps can be operated with multiple RF drive signals [45, 74, 75, 76]. Therefore, we generalize the theoretical description to the multi-drive case, where interference between drives introduces additional cross-terms. For the derivation of the model, we can follow the same steps as in the previous case. We begin with the squared magnitude of the total electric field including the injected noise, perform the time average, and then identify the resulting force components. From these expressions, the corresponding heating rates follow directly, now containing both self-terms (from individual drives) and cross-terms (arising from interference between different driving signals).

The electric field can be written, in full generality, as the sum of k electric fields, each of them corresponding to a different RF signal with the same form of Equation 3.14:

$$\begin{aligned}\vec{E}_{\mathrm{RF}}(x,y,z,t) &= \sum_k \vec{E}^{(k)}_{\mathrm{RF}}(x,y,z,t) = \\ &= \sum_k \vec{E}^{(k)}_0(x,y,z)\left[\cos(\Omega_{\mathrm{RF}}t) + \chi_N^{(k)}\cos\left((\Omega_{\mathrm{RF}}+\omega)t + \varphi^{(k)}(t)\right)\right], \end{aligned} \tag{3.23}$$

where $\varphi^{(k)}(t)$ represents a stochastic phase fluctuation, which accounts for correlations between the noise contributions of different RF sources. Possible static phase offsets of the individual RF tones are omitted here, since they average out in the subsequent time-averaging procedure and do not contribute to the effective noise force.

The total electric field is composed of several contributions, thus, when computing the squared magnitude, two types of terms naturally appear:

1. The first type comes from squaring the individual contributions. Since each contribution is sum of two sub-terms, squaring it also generates internal cross-products between its sub-terms.

2. The second type consists of cross-products between different contributions. These involve two indices, k and l, corresponding to scalar products between distinct electric field components, $\vec{E}_k \cdot \vec{E}_l$ with $k \neq l$.

For the rest of this chapter, we will refer to the first type as *diagonal terms* and to the second type as *off-diagonal terms.* Explicitly, the squared magnitude of the electric field becomes:

$$\begin{aligned}
\left|\vec{E}_{\mathrm{RF}}\right|^2 &= \left|\sum_k \vec{E}_0^{(k)}\left[\cos(\Omega_{\mathrm{RF}}t) + \chi_N^{(k)}\cos\left((\Omega_{\mathrm{RF}}+\omega)t + \varphi^{(k)}(t)\right)\right]\right|^2 = \\
&= \sum_k \left|\vec{E}_0^{(k)}\right|^2 \cos^2(\Omega_{\mathrm{RF}}t) + \sum_{k\neq l}\left(\vec{E}_0^{(k)}\cdot\vec{E}_0^{(l)}\right)\cos^2(\Omega_{\mathrm{RF}}t) + \\
&\quad + 2\sum_k \left|\vec{E}_0^{(k)}\right|^2 \chi_N^{(k)}\cos\left((\Omega_{\mathrm{RF}}+\omega)t + \varphi^{(k)}(t)\right)\cos(\Omega_{\mathrm{RF}}t) + \\
&\quad + \sum_{k\neq l}\left(\vec{E}_0^{(k)}\cdot\vec{E}_0^{(l)}\right)\chi_N^{(k)}\cos(\Omega_{\mathrm{RF}}t)\cos\left((\Omega_{\mathrm{RF}}+\omega)t + \varphi^{(k)}(t)\right) + \\
&\quad + \sum_{k\neq l}\left(\vec{E}_0^{(k)}\cdot\vec{E}_0^{(l)}\right)\chi_N^{(l)}\cos(\Omega_{\mathrm{RF}}t)\cos\left((\Omega_{\mathrm{RF}}+\omega)t + \varphi^{(l)}(t)\right) + \\
&\quad + \mathcal{O}(\chi_N^2)\,,
\end{aligned} \tag{3.24}$$

where, for readability, terms quadratic in the noise amplitudes, $\mathcal{O}(\chi_N^2)$, are not written explicitly. Compared to the single-drive case, additional linear contributions now appear due to interference between different RF fields, which can either enhance or suppress the ion's heating. Note that the last two sums are identical up to a dummy–index exchange, since the dot product is symmetric; therefore, they can be combined into a single double sum together with the diagonal term. In particular, since $\left|\vec{E}_0^{(k)}\right|^2 = \vec{E}_0^{(k)}\cdot\vec{E}_0^{(k)}$, the complete expression may be written as a sum over all pairs (k,l):

$$\begin{aligned}
\left|\vec{E}_{\mathrm{RF}}\right|^2 &= \sum_k \left|\vec{E}_0^{(k)}\right|^2 \cos^2(\Omega_{\mathrm{RF}}t) + \sum_{k\neq l}\left(\vec{E}_0^{(k)}\cdot\vec{E}_0^{(l)}\right)\cos^2(\Omega_{\mathrm{RF}}t) + \\
&\quad + 2\sum_{k,l}\left(\vec{E}_0^{(k)}\cdot\vec{E}_0^{(l)}\right)\chi_N^{(k)}\cos(\Omega_{\mathrm{RF}}t)\cos\left((\Omega_{\mathrm{RF}}+\omega)t + \varphi^{(k)}(t)\right) + \mathcal{O}(\chi_N^2),
\end{aligned} \tag{3.25}$$

Again, Werner's formula can be used to expand the products of cosines in the cross

term; as previously explained, the fast oscillating terms coming from the expansion vanish under the RF-cycle average, while the slow components are assumed constant and can be taken outside the integral.

Thus, after applying the time average operation (up to the second order in the noise), we get:

$$\left\langle \left|\vec{E}_{\mathrm{RF}}\right|^2 \right\rangle \approx \frac{1}{2}\left|\sum_k \vec{E}_0^{(k)}\right|^2 + \sum_{k,l}\left(\vec{E}_0^{(k)}\cdot\vec{E}_0^{(l)}\right)\chi_N^{(k)}\cos\left(\omega t+\varphi^{(k)}(t)\right).$$

Again, the first term corresponds to the noiseless pseudopotential arising in the absence of noise, while the other terms represent the noise-induced modulations.

During the time-averaging procedure, each $\varphi^{(k)}(t)$ is assumed to fluctuate randomly with zero mean, so its contribution averages out and can be neglected. Whether the cross-terms contribute to the effective field, however, depends on the correlation properties of the stochastic phases $\varphi^{(k)}(t)$:

- If the noise is *completely correlated*, i.e. $\varphi^{(k)}(t) = \varphi^{(l)}(t)\ \forall k,l$, and the noise amplitudes are identical (a reasonable assumption for white noise affecting similar RF signals), then the expression for Φ_N at the ion position is:

$$\Phi_N = \frac{Q^2}{4m\Omega_{\mathrm{RF}}^2}\left|\sum_k \vec{E}_0^{(k)}\right|^2 \chi_N\cos\left(\omega t+\varphi(t)\right) \tag{3.26}$$

  The noise force comes from Equation 3.13:

$$\vec{F}_N = -\frac{2Q^2}{4m\Omega_{\mathrm{RF}}^2}\left(\vec{\nabla}\left|\sum_k \vec{E}_0^{(k)}\right|^2\right)\chi_N\cos\left(\omega t+\varphi(t)\right) \tag{3.27}$$

  This case is similar to the single drive result, because only diagonal terms contribute (cfr. Equation 3.19). Thus, we can use the relation $\vec{\nabla}\left|\vec{A}\right|^2 = 2\left|\vec{A}\right|\left(\vec{\nabla}\left|\vec{A}\right|\right)$, and, since $\vec{E} = \sum_k \vec{E}_0^{(k)} = 0$ at the ion position, we always obtain $\vec{F}_N = 0$. Therefore, for completely correlated noise, the ion experiences no net force in the pseudopotential approximation, and the heating rate is effectively zero.

- If the noise is *uncorrelated*, i.e. if the $\varphi^{(k)}(t)$ are independent random processes, the noise force contribution becomes:

$$\vec{F}_N = \frac{Q^2}{4m\Omega_{\mathrm{RF}}^2}\sum_k\left[\left(\sum_l \vec{\nabla}\left[\vec{E}_0^{(k)}\cdot\vec{E}_0^{(l)}\right]\right)\chi_N^{(k)}\cos\left(\omega t+\varphi^{(k)}(t)\right)\right] \tag{3.28}$$

In this case, the noise terms carry independent random phases, which prevents the oscillations from adding coherently; even if the signals share the same modulation amplitude, a residual net force is present at the ion position. The reason is that the relative dephasing introduced by the $\varphi^{(k)}(t)$ prevents the cancellation that occurred in the correlated case. In particular, we can no longer rewrite the expression for the gradient of the scalar product, since the off-diagonal contributions are now present. As a result, the ion experiences a fluctuating force, differently from the previous case, even if confined in the RF null.

Having established that a surviving net force arises only in the case of uncorrelated noise, we can now derive an expression for the heating rate, combining Equation 3.10 and Equation 3.11. To this end, we require the squared magnitude of the vector $\vec{F}_0$ as well as the corresponding spectral force density, expressed by Equation 3.22. Starting from the expression of the noise force (Equation 3.28), we note that when expanding the square, cross terms appear in the form $\chi_N^{(k)}\chi_N^{(l)}$ with $k \neq l$. However, these cross terms vanish upon averaging. Specifically, at $t = 0$ they contain products of the form $\cos\varphi^{(k)}(0)\,\cos\varphi^{(l)}(0)$ where the phases $\varphi^{(k)}(0)$ are randomly distributed. Since the ensemble average of products of independent random phases factorizes, we obtain:

$$\left\langle \cos\varphi^{(k)}(0)\,\cos\varphi^{(l)}(0)\right\rangle = \left\langle\cos\varphi^{(k)}(0)\right\rangle\left\langle\cos\varphi^{(l)}(0)\right\rangle = 0, \tag{3.29}$$

which implies that all cross terms vanish. Consequently, only the diagonal contributions in $\left(\chi_N^{(k)}\right)^2$ survive in the squared magnitude and the total heating rate is:

$$\dot{\bar{n}} = \frac{Q^4}{16m^3\Omega_{\mathrm{RF}}^4\hbar\omega}\sum_k\left[\left(\sum_l \vec{\nabla}\left[\vec{E}_0^{(k)}\cdot\vec{E}_0^{(l)}\right]\right)^2 \frac{S_{V_N}^{(k)}(\Omega_{\mathrm{RF}}+\omega)}{\left(V_0^{(k)}\right)^2}\right]. \tag{3.30}$$

From this expression it's clear that the total heating rate for uncorrelated noise depends both on the spectral properties of the noise sources and on the geometric factors of the trap fields. This dependency is the same that we have obtained previously for the single RF drive case, expressed in Equation 3.20. Again, the double summation encodes the electrode geometry: for each noisy RF signal k, the effective contribution to heating is weighted by the gradients of the cross-terms $\vec{E}_0^{(k)}\cdot\vec{E}_0^{(l)}$ with all other signals l. In contrast to the correlated case, where only diagonal terms survive and cancellation can occur, here the off-diagonal contributions remain, and their squared sum ensures that incoherent fluctuations always produce a nonzero heating rate. Finally, normalizing the voltage noise spectral density by the squared drive amplitude ensures that the heating

rate reflects the relative noise on each electrode, rather than its absolute drive amplitude.

3.2.2 Phase noise model

Similarly to what we did for the amplitude-noise case, we begin the analysis of the effect of phase noise by considering a single RF drive signal, and then we present the generalization to a trap driven by multiple RF signals.

In this case, the modulation acts directly on the phase of the drive signal. Differently from the amplitude-noise case, the perturbation is embedded within the oscillation itself, rather than appearing as an additional sideband. As a consequence, it is sufficient to describe the modulation as a slow variation of the phase, which can resonantly couple to the ion's secular motion and lead to excitation. The resulting electric field takes the form:

$$\vec{E}_{\mathrm{RF}}(x,y,z,t) = \vec{E}_0(x,y,z)\left[\cos\left(\Omega_{\mathrm{RF}}t + \theta_N(t)\right)\right], \tag{3.31}$$

where $\theta_N(t) \equiv \chi_\theta \cos(\omega_N t)$ represents the phase noise with modulation amplitude $\chi_\theta \ll 1$ and frequency $|\omega_N| \ll \Omega_{\mathrm{RF}}$. Since $\chi_\theta \ll 1$, we can first expand the cosine of a sum using the standard identity $\cos(A+B) = \cos A \cos B - \sin A \sin B$, and then apply the Taylor expansion for small $\theta_N(t)$. Up to the second order in $\theta_N(t)$, this gives:

$$\vec{E}_{\mathrm{RF}}(x,y,z,t) \approx \vec{E}_0(x,y,z)\left[\cos(\Omega_{\mathrm{RF}}t)\left(1 - \frac{\theta_N^2(t)}{2}\right) - \sin(\Omega_{\mathrm{RF}}t)\theta_N(t)\right]. \tag{3.32}$$

From this, we compute the squared magnitude of the electric field, keeping only up to the second order in $\theta_N(t)$:

$$\begin{aligned}\left|\vec{E}_{\mathrm{RF}}\right|^2 = \left|\vec{E}_0\right|^2 &\left[\cos^2(\Omega_{\mathrm{RF}}t)\left(1-\theta_N^2(t)\right) + \sin^2(\Omega_{\mathrm{RF}}t)\theta_N^2(t)\right] + \\ &- 2\left|\vec{E}_0\right|^2 \cos(\Omega_{\mathrm{RF}}t)\sin(\Omega_{\mathrm{RF}}t)\theta_N(t)\end{aligned} \tag{3.33}$$

Since the noise frequency is much smaller than the RF drive, $|\omega| \ll \Omega_{\mathrm{RF}}$, the terms proportional to $\theta_N(t)$ and $\theta_N^2(t)$ can be considered constant over a single RF cycle and can be taken out of the time average. First of all, noting that $\langle\cos^2(\Omega_{\mathrm{RF}}t)\rangle_T = \langle\sin^2(\Omega_{\mathrm{RF}}t)\rangle_T = 1/2$, the terms proportional to $\theta_N^2(t)$ cancel out. Moreover, the last term vanishes, because $\langle\cos(\Omega_{\mathrm{RF}}t)\sin(\Omega_{\mathrm{RF}}t)\rangle_T = 0$. Thus, the pseudopotential reduces to:

$$\Phi_{\mathrm{PS}} = \frac{Q^2}{4m\Omega_{\mathrm{RF}}^2}\left\langle\left|\vec{E}_0\right|^2\right\rangle = \Phi_0\,, \tag{3.34}$$

which coincides with the usual expression for the pseudopotential and shows no additional contribution due to phase noise. This means that no net noise force acts on the ion and, consequently, no heating is observed in this case, leading to the conclusion that phase noise does not lead to motional excitation when considering a single-drive configuration.

We now generalize the derivation to the multi-drive case, in order to examine whether heating arises from relative phase differences. The procedure follows the same steps as before, where now the total electric field is written as the sum of k contributions, each carrying its own phase modulation $\theta_N^{(k)}(t)$:

$$\vec{E}_{\mathrm{RF}}(x,y,z,t) = \sum_k \vec{E}_{\mathrm{RF}}^{(k)}(x,y,z,t) = \sum_k \vec{E}_0^{(k)}(x,y,z)\left[\cos\left(\Omega_{\mathrm{RF}}t + \theta_N^{(k)}(t)\right)\right]. \tag{3.35}$$

In contrast to the amplitude-noise case, there is no need to introduce an additional random modulation $\varphi^{(k)}(t)$ to each signal in order to express correlations between different noise sources. The phase terms $\theta_N^{(k)}(t)$ already incorporate these correlations: if all phases are identical, the noise is fully correlated across the drives, whereas deviations between them directly quantify the degree of correlation. Introducing an extra stochastic variable would therefore be redundant. Again, we can expand the cosine of the sum and use a Taylor expansion to simplify the expression of the electric field:

$$\vec{E}_{\mathrm{RF}} \approx \sum_k \vec{E}_0^{(k)}\left[\cos(\Omega_{\mathrm{RF}}t)\left(1 - \frac{\left(\theta_N^{(k)}(t)\right)^2}{2}\right) - \sin(\Omega_{\mathrm{RF}}t)\theta_N^{(k)}(t)\right]. \tag{3.36}$$

Reorganizing the terms, the square amplitude of this vector becomes:

$$\begin{aligned}\left|\vec{E}_{\mathrm{RF}}\right|^2 = &\sum_{k,l}\left(\vec{E}_0^{(k)}\cdot\vec{E}_0^{(l)}\right)\left[\cos^2(\Omega_{\mathrm{RF}}t)\left(1 - \frac{\left(\theta_N^{(k)}(t)\right)^2 + \left(\theta_N^{(l)}(t)\right)^2}{2}\right)\right] + \\ &+ \sum_{k,l}\left(\vec{E}_0^{(k)}\cdot\vec{E}_0^{(l)}\right)\left[\sin^2(\Omega_{\mathrm{RF}}t)\left(\theta_N^{(k)}(t)\theta_N^{(l)}(t)\right)\right] + \\ &- \sum_{k,l}\left(\vec{E}_0^{(k)}\cdot\vec{E}_0^{(l)}\right)\cos(\Omega_{\mathrm{RF}}t)\sin(\Omega_{\mathrm{RF}}t)\left(\theta_N^{(k)}(t) + \theta_N^{(l)}(t)\right),\end{aligned} \tag{3.37}$$

where again only terms up to the second order in the noise amplitude have been kept. Similarly at the single-drive case, also now the phase noise component can be considered constant as $|\omega| \ll \Omega_{\mathrm{RF}}$ and can be taken out of the integral when performing the time

average. The pseudopotential thus becomes:

$$\Phi_{\mathrm{PS}} = \frac{Q^2}{4m\Omega_{\mathrm{RF}}^2}\left|\sum_k \vec{E}_0^{(k)}\right|^2 - \frac{Q^2}{2m\Omega_{\mathrm{RF}}^2}\sum_{k\neq l}\left(\vec{E}_0^{(k)}\cdot\vec{E}_0^{(l)}\right)\left[\frac{\theta_N^{(k)}(t)-\theta_N^{(l)}(t)}{2}\right]^2 = \Phi_0 + \Phi_N \quad (3.38)$$

where only the terms quadratic in the noise amplitude survive the time average. As we can see, unlike the single-drive case, a term now survives that depends on the phase difference between the different signals. This term has no diagonal component and, therefore, does not appear for a single drive. Since the noise contribution to the pseudopotential can only excite the ion motion if resonant with the secular frequencies, we can, without loss of generality, replace each phase difference with the following expression:

$$\alpha_{kl}(t) = \chi_{kl}\cos(\omega_N t) = \frac{\theta_N^{(k)}(t)-\theta_N^{(l)}(t)}{2}\,, \quad (3.39)$$

where χ_{kl} represent the amplitude of the combination of the noise on the two signals and ω_N is the generic noise frequency.

If we substitute Equation 3.39 in the expression for Φ_N, we see that there is an important difference with the result obtained in the amplitude-noise case. This time, Φ_N depends on the square of a cosine, which implies that the resulting contribution oscillates at twice the noise frequency. Explicitly we obtain:

$$\Phi_N = -\frac{Q^2}{4m\Omega_{\mathrm{RF}}^2}\sum_{k\neq l}\left(\vec{E}_0^{(k)}\cdot\vec{E}_0^{(l)}\right)\chi_{kl}^2 - \frac{Q^2}{4m\Omega_{\mathrm{RF}}^2}\sum_{k\neq l}\left(\vec{E}_0^{(k)}\cdot\vec{E}_0^{(l)}\right)\chi_{kl}^2\cos(2\omega_N t) \quad (3.40)$$

The noise component of the pseudopotential is the sum of two contributions: a constant energy variation and a time-dependent term. The static part of the noise alters the pseudopotential by adding small off-diagonal elements. As a result, whereas Φ_0 in Equation 3.38 initially included only diagonal contributions, the corrected pseudopotential now exhibits couplings between different spatial directions. Conversely, the time-dependent component is responsible for the net noise force acting on the ion. It is important to emphasize that the noise frequency ω_N must equal half the secular frequency in order to be resonant with the ion motion, differently from what we obtained for the amplitude-noise case.

Using Equation 3.13, the noise force is:

$$\vec{F}_N = -\frac{Q^2}{4m\Omega_{\mathrm{RF}}^2}\sum_{k\neq l}\left(\vec{\nabla}\left[\vec{E}_0^{(k)}\cdot\vec{E}_0^{(l)}\right]\right)\chi_{kl}^2\cos(2\omega_N t). \quad (3.41)$$

The phase noise spectral density comes from Equation 3.5:

$$\begin{aligned} S_{\Theta_N}(\Omega_{\mathrm{RF}} \pm \omega) &= \left\langle \chi_\theta^2 \right\rangle \int_{-\infty}^{+\infty} \langle \cos(2\omega_N t) \cos(2\omega_N (t+\tau)) \rangle_t \; e^{-i(\Omega_{\mathrm{RF}} \pm \omega)\tau} d\tau \\ &= \frac{1}{2} \left\langle \chi_\theta^2 \right\rangle \int_{-\infty}^{+\infty} \cos(2\omega_N \tau) \; e^{-i(\Omega_{\mathrm{RF}} \pm \omega)\tau} d\tau \\ &= \frac{\pi}{2} \left\langle \chi_\theta^2 \right\rangle [\delta(\omega - 2\omega_N) + \delta(\omega + 2\omega_N)] \end{aligned} \tag{3.42}$$

and then can be related to the noise force spectral density thanks to Equation 3.12.

Combining Equation 3.42 with Equation 3.10, we obtain an expression for the heating rate in the case of phase noise:

$$\dot{\bar{n}} = \frac{Q^4}{16 m^3 \Omega_{\mathrm{RF}}^4 \hbar \omega} \left[\left(\sum_{k \neq l} \vec{\nabla} \left[\vec{E}_0^{(k)} \cdot \vec{E}_0^{(l)} \right] \right)^2 S_{\Theta_N}^{(k)} \left(\Omega_{\mathrm{RF}} + \frac{\omega}{2} \right) \right] \tag{3.43}$$

Thus, also in the case of phase noise, the total heating rate depends both on the spectral properties of the noise sources and on the geometric factors of the trap fields. However, these dependencies differ from the amplitude noise case. First, the diagonal term in the double summation is absent, since we have shown that phase noise on a single RF drive does not produce any heating. Second, the phase-noise spectral density does not require normalization by the drive amplitude, because it is already expressed as a relative fluctuation of the RF signal itself.

In the following chapters, we introduce the *Quantum Charge-Coupled Device* (QCCD) and the *Quantum Spring Array* (QSA) architectures for surface-electrode ion traps and discuss their key operational elements. For each platform, we identify where the heating models developed in this chapter become relevant. For the QCCD architecture, the analysis focuses on X-junctions, while for the QSA architecture, the study concentrates on the RF-transport mechanisms.

4

RF technical noise in the QCCD architecture

In this chapter, we investigate the effects of RF technical noise on the X-junction trap in the QCCD architecture. The first section introduces the fundamental concepts of the QCCD architecture. The second section focuses on the characterization of the X-junction: after a brief description, the primary sources of heating in the junction are analyzed. The third section covers the optimization of the X-junction, carried out through a minimization procedure aimed at reducing the geometrical contribution to the heating. Finally, the last section presents heating-rate simulations obtained by applying the noise models developed in Chapter 3 to the X-junction, and compares the results between the non-optimized and optimized configurations.

4.1 The QCCD architecture

The QCCD architecture [34] represents a prominent blueprint for scalable trapped-ion quantum computers. It was conceived to overcome the physical limitations of operating with long ion chains by distributing computation across multiple interconnected trapping zones. In a QCCD processor, individual ions are confined, manipulated, and measured in distinct regions of a single chip, while being physically *shuttled*, i.e. transported, between zones as needed. A typical layout includes *storage zones* for idle qubits, *logic zones* for entangling operations, *cooling and detection zones* for state initialization and readout, and a network of junctions and transport channels that interconnect them, as illustrated in Figure 4.1.

This spatial separation of functions allows each region to be optimized independently and the modular structure enables parallel operation: gate operations, cooling and read-

out can be performed in different zones simultaneously, and this constitutes an essential feature for scaling to larger qubit numbers. By shuttling ions rather than relying on static, long-range coupling, the QCCD achieves reconfigurable connectivity between qubits without increasing the gate complexity. From an engineering perspective, the implementation of the QCCD in planar surface traps relies on precise control of segmented DC electrodes to generate the transport waveforms needed for moving ions between neighbouring trap zones. Different strategies for producing these waveforms and managing ion motion are discussed in Section A.2.

A key design requirement of the QCCD architecture is the presence of *junctions*, building blocks that allow ions to change direction during transport and enable full two-dimensional connectivity between trapping zones. These junctions are essential for large-scale QCCD processors, as they connect linear channels into a navigable network and permit flexible routing of qubits across the device. For example, this structure allows the controlled transport of qubits from dedicated memory zones to interaction regions where quantum logic operations are performed [34]. The specific junction ge-

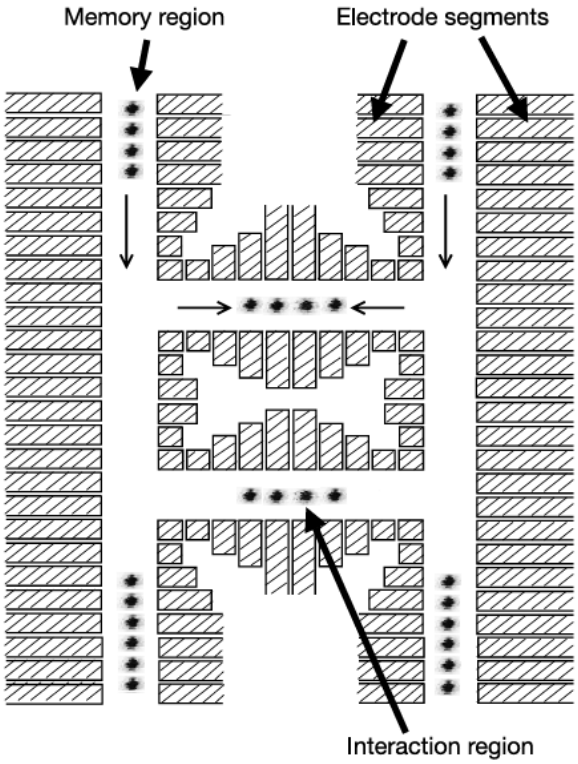

FIGURE 4.1: Diagram of the QCCD architecture. Ions are stored in the memory region and moved to the interaction region for logic operations. Thin arrows show transport and confinement along the local trap axis [34].

ometry has a strong impact on transport efficiency, motional heating, and sensitivity to RF noise. This is because junctions create regions in which RF fields from different branch directions overlap and interfere, producing complex field gradients and potential barriers that must be carefully controlled during shuttling.

Several junction geometries have been explored in the literature, developed for different architectural needs, including T-junctions [77], Y-junctions [60, 78], and X-junctions [41, 79, 80, 81]. The focus of this chapter is on the X-junction, which constitute the predominant junction configuration adopted in current experimental implementations of the QCCD architecture, due to its geometric symmetry and suitability for scalable layouts. In the following section, we introduce the X-junction geometry in more detail, presenting its properties and how they are related to the motional heating of the ion.

4.2 X-junction characterization

We begin by analyzing the physical characteristics of the X-junction, which, compared to linear trap sections, presents significant design challenges due to the superposition of RF fields from intersecting channels. This overlap distorts the ideal quadrupolar confinement and gives rise to a pseudopotential barrier near the junction center [41, 60, 79, 80]. As a result, the effective trapping strength decreases in this region, making it more difficult to maintain stable confinement during ion transport. These distortions impede smooth ion transport across the junction and can lead to excess micromotion as the RF field is no longer vanishing in the pseudopotential minimum. Furthermore, the presence of local potential gradients increases the risk of motional excitation and heating, as discussed in Chapter 3, ultimately degrading the coherence and fidelity of quantum operations. In surface-electrode traps, these effects are further amplified by the lack of mirror or rotational symmetry perpendicular to the trap surface. This geometric constraint naturally reduces the trap depth [61] and enhances the ion's sensitivity to electric field noise originating from the electrodes [41]. Consequently, precise modeling and optimization of the electric field near the junction center are essential. In this section, we focus on analyzing the electric field and the pseudopotential across the entire junction, identifying the origins of the weakened confinement and pseudopotential barrier. This understanding forms the basis for the optimization of the junction geometry aimed at minimizing micromotion and motional heating along the junction, thereby enabling reliable, high-fidelity ion transport between different regions of the QCCD architecture.

4.2.1 X-JUNCTION GEOMETRY AND DESIGN

An X-junction is formed by the intersection of two linear ion-trap channels, creating a four-way crossing. Each arm of the junction can be conceptually divided into two regions. The *linear region* behaves like a standard linear trap, where the RF and DC fields produce near-ideal quadrupolar confinement, allowing ions to be transported with minimal heating. The *center region*, in contrast, is where the RF fields from the intersecting channels overlap and interfere, generating distortions of the ideal confinement and potential barriers that dominate the ion dynamics.

In an ideal X-junction, the RF electric field would vanish along the path identified by the pseudopotential minimum through the junction, ensuring minimal RF-driven micromotion and motional heating during transport. By symmetry, the RF field vanishes exactly at the geometric center, but deviations in the surrounding region are critical for transport and motional heating. To quantify and understand these deviations, it is convenient to analyze the local structure of the electric field using a *multipole expansion* around the junction center [82]. In this framework, the electric field near the intersection is expressed as a series of terms with increasing order, each characterized by a specific symmetry and spatial dependence:

$$E_i\left(\vec{r}\right) = d_i + \sum_j q_{i,j} r_j + \frac{1}{2}\sum_j \sum_k h_{i,j,k} r_j r_k + \mathcal{O}(|\vec{r}|^3), \tag{4.1}$$

where d_i is the dipole term, and the tensors $q_{i,j} = \partial_j E_i$ and $h_{i,j,k} = \partial_j \partial_k E_i$ correspond to the quadrupole and hexapole components of the field, respectively.

We want to determine which multipole term dominates the field in the central region of the junction; since a dipole field is spatially uniform and cannot confine charged particles, at least a quadrupole term is required to achieve trapping. The order of the lowest non-vanishing multipole term at the trap center is particularly significant for two main reasons:

- **Confinement strength and heating**: For a given ion–surface distance, lower-order multipoles provide stronger confinement [82]. Maintaining strong confinement at larger ion–surface distances is desirable, since ion heating from microscopic surface noise increases rapidly as the ion approaches the electrodes [40], as discussed in Section 3.1.

- **Feasibility of junction design**: The RF electric field at the junction center vanishes if the dipole term d_i is zero. The shape of the potential near the center, which determines ion confinement and transport, is then set by the lowest-order multipole

terms allowed by the junction symmetry (quadrupole, hexapole, etc.). Designing a junction such that confinement can be realized using only lower-order terms simplifies the electrode geometry, reduces the number of symmetry constraints, and results in a structure that is more robust to fabrication imperfections [82].

It has been shown that, for an ideal X-junction in which the RF electric field vanishes along both intersecting ion transport paths while still providing confinement, the dipole and quadrupole components of the field must be zero [82]. Thus, the lowest-order multipole that can survive at the intersection is a hexapole, which is higher order than the quadrupole that dominates the field in a standard linear trap (expressed in Equation 2.7). However, an ion located at the center of a pure hexapole potential experiences zero harmonic motional frequency [83], meaning that there is no quadratic restoring force to produce standard harmonic oscillations. In other words, the ion cannot be confined in the usual sense of a harmonic trap, and higher-order confinement is both weaker and highly nonlinear. This limitation illustrates why perfectly nulling all low-order multipoles is not practical: it is not possible to simultaneously achieve strong harmonic confinement and a perfectly zero RF field along all transport paths.

For this reason, we will consider a more realistic junction design, where a non-zero RF quadrupole is typically retained at the intersection to provide sufficient confinement. In this case, the RF field is inevitably *non-zero* along the ion paths approaching or leaving the center. Because the axial translational symmetry is broken in the central region, this RF field is inhomogeneous along the trap axis, producing a *pseudopotential barrier* [41, 80]. This inhomogeneous RF field is particularly relevant for technical noise: fluctuations in the RF drive, indeed, can couple directly to the ion's motion along these paths, leading to additional motional heating during transport, as discussed in Chapter 3.

To mitigate these effects and provide a stable framework for transport through the junction, the RF electrode geometry can be optimized to reduce the height of the pseudopotential barrier, thereby minimizing the heating experienced by ions when shuttled through the junction. This will be the focus of Section 4.3.

Before optimizing the junction itself, the geometry of the RF electrodes in the linear trap sections, which are dominated by the quadrupole potential, can be optimized to maximize the pseudopotential curvature at the chosen ion height h, which is directly related to the trapping strength (see Equation 2.21) [62, 80]. Strong confinement in these linear sections provides a well-characterized baseline for integrating the junction with the surrounding trap structure, which is the ultimate objective for the QCCD architecture.

For this reason, we first consider an X-junction whose geometrical parameters match those of the linear sections (as shown in Figure 4.2, where only the RF electrodes are considered). In these regions, the primary adjustable design parameters, without changing the electrode shapes, are the gap between the RF electrodes, w_G, and their width, w_{RF}. By maximizing the pseudopotential curvature at a given ion height h, it is possible to derive a relationship between these geometrical parameters and the desired ion–surface distance [80]:

$$w_G = 0.83h \qquad w_{RF} = \frac{4h^2 - w_G^2}{2w_G} \approx 1.99h \tag{4.2}$$

The different ion–surface distances, together with the corresponding geometrical parameters studied in this work, are reported in Table 4.1.

In the following section, we present simulations of the non-optimized X-junction, analyzing its key physical quantities and the resulting ion dynamics during shuttling through the trap. Before proceeding, it is necessary to specify the fixed parameters

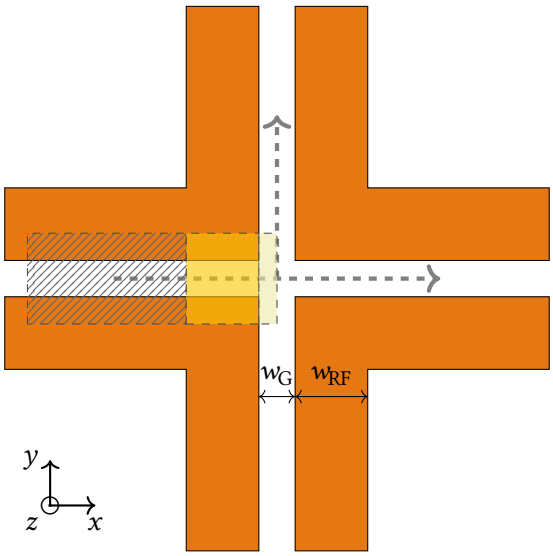

FIGURE 4.2: Schematic representation of an X-junction formed by the intersection of two linear traps. The orange regions indicate the RF electrodes, each of width w_{RF}. The central cross-shaped gap of width w_G, together with the surrounding square regions outside the RF electrodes, defines the space reserved for DC electrodes. As in linear traps, the ion is confined in the region between two RF electrodes, and at the junction it can either be shuttled straight through or redirected by 90°. The regions enclosed by dashed borders on the left arm of the junction indicate the area used for the simulations. The first area still resembles a linear trap, while the two shaded yellow regions correspond to the portions of the junction arm associated with the perpendicular RF electrodes (darker area) and with the central opening of the junction (lighter area).

Ion height h	RF width w_{RF}	Gap width w_G
50 µm	99.5 µm	41.5 µm
75 µm	149.25 µm	62.25 µm
100 µm	199 µm	83 µm

Table 4.1: Ion–trap distances with corresponding RF electrode widths and gaps, optimized to maximize the pseudopotential curvature at the chosen ion height.

used in these calculations. We consider a single trapped ion, assumed to be $^{40}Ca^{+}$ with mass $m = 40\,\mathrm{u}$ and charge $Q = +e$. Unless stated otherwise, all simulations presented in this section assume a fixed ion–surface distance of $h = 100\,\mu\mathrm{m}$, with RF electrode dimensions taken from Table 4.1. Simulations with different ion–surface distances are shown in Section A.1. The ion is confined by setting a Mathieu parameter $q = 0.4$ and a RF drive frequency of $\Omega_{RF} = 20\,\mathrm{MHz}$. Using the general relation between the Mathieu parameters and the applied voltages (Equation 2.10), the corresponding RF voltage amplitude is given by:

$$V_{RF} = q\,\frac{mh^2\Omega_{RF}^2}{2Q}\,, \tag{4.3}$$

and for the chosen parameters, yields $V_{RF} \approx 42.2\,\mathrm{V}$.

Finally, the junction should be simulated assuming infinitely long arms, so that boundary effects do not influence the central region. To approximate this condition, the arms have been chosen to be more than an order of magnitude longer than the typical width of the RF electrodes (approximately $15 \times w_{RF}$), which has been shown to be enough to preserve the pseudopotential landscape near the junction center.

4.2.2 Pseudopotential analysis

Building upon the geometric characterization presented in the previous subsection, we now analyze the pseudopotential generated by the RF electrodes in the X-junction. This study provides quantitative insight into the trapping efficiency and identifies the regions that are most critical for ensuring smooth ion transport through the junction. The main quantities of interest include the pseudopotential at the ion position, its gradient, the equilibrium ion position (corresponding to the pseudopotential minimum), as well as the secular frequencies and their variance.

The analysis consists of simulating ion shuttling along one arm of the junction toward its center, i.e. studying the transport path in which the ion is always placed at the corresponding pseudopotential minimum for each axial position [80]. Due to the symmetry

of the structure, studying a single arm is sufficient to represent both a sharp 90° turn and transport through the full junction. The shuttling trajectory is modelled by computing the pseudopotential minimum over a grid of points along the arm, using a spatial step $\Delta x = 0.5\,\mu\text{m}$. The pseudopotential is evaluated at each position, together with its gradient and the z-coordinate of the minimum, while maintaining the reference frame defined in Figure 4.2. The resulting configuration is illustrated in Figure 4.3, where the two shaded yellow regions indicate the portions of the junction arm corresponding to the perpendicular RF electrodes (darker area) and the central opening of the junction (lighter area). These regions mark the transition from the linear section of the trap to the central crossing, where the influence of the additional RF pair becomes significant,

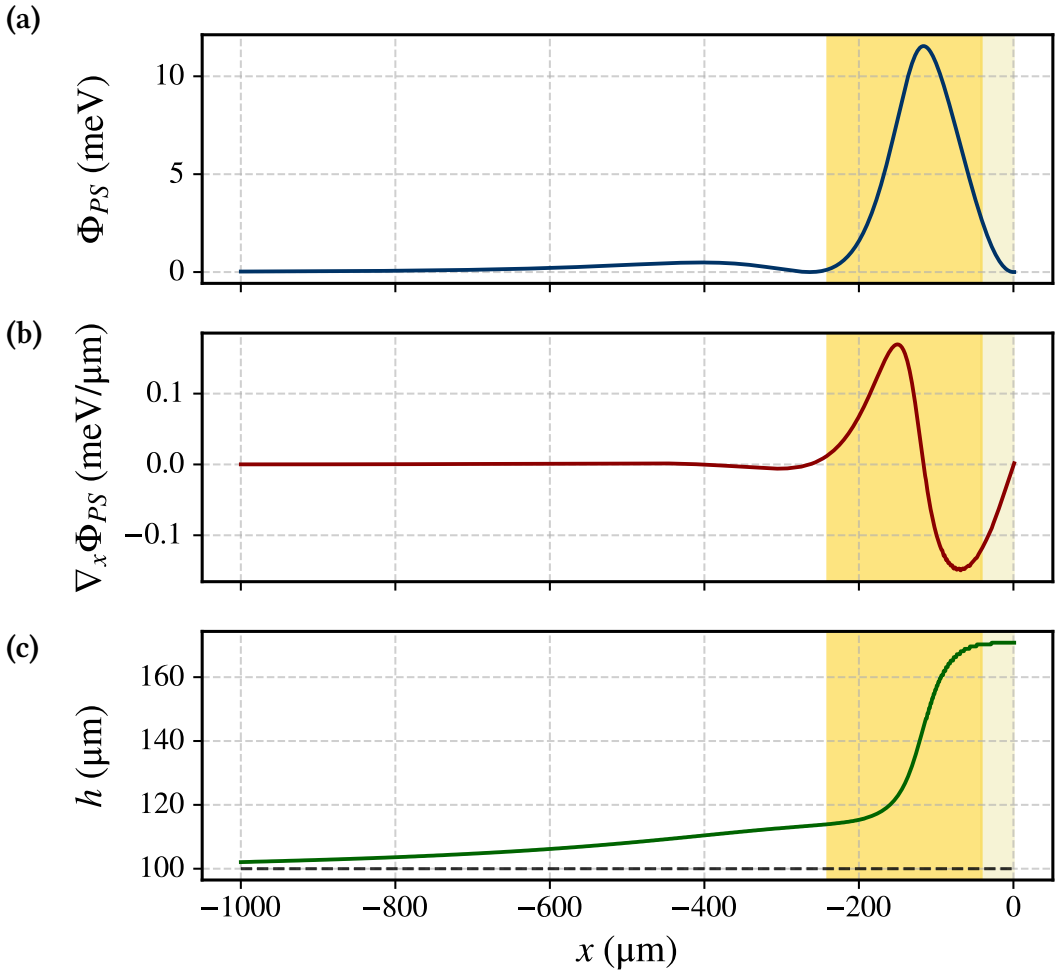

Figure 4.3: Pseudopotential analysis along one arm of the non-optimized X-junction. (a) Pseudopotential profile $\Phi_{\text{PS}}(x)$, showing the pseduopotential barrier peaking at $x \approx 95\,\mu\text{m}$ with values exceeding 10 meV. (b) Corresponding gradient, exhibiting behavior consistent with the pseudopotential profile. (c) z-coordinate of the pseudopotential minimum, compared to the target ion–surface distance $h = 100\,\mu\text{m}$.

as shown in Figure 4.2.

The pseudopotential barrier described in Subsection 4.2.1 is clearly visible in panel (a). The barrier peak occurs at approximately $x \approx 95\,\mu\text{m}$, near the midpoint of the perpendicular RF electrode, and its maximum value is of 11.56 meV. It is worth noting that, already within the linear section of the junction (up to $x \approx -250\,\mu\text{m}$) a slight increase in the pseudopotential is observed. At the junction center, the pseudopotential returns to very low values (lower than 0.01 meV).

Panel (b) shows the gradient of the pseudopotential along the x-axis, obtained by directly differentiating the pseudopotential. This quantity is of particular importance for this work: as shown in Chapter 3, it enters directly into the expression for the heating rate for RF technical noise. Consequently, in the following sections, we will explore strategies to reduce the pseudopotential barrier in order to minimize this gradient and thereby mitigate ion heating during transport.

Panel (c), on the other hand, presents the z-coordinate of the pseudopotential minimum, i.e. the ion transport path, compared to the target height $h = 100\,\mu\text{m}$. It can be observed that, even far from the junction center, the ion–surface distance does not exactly coincide with h, indicating that the junction's influence extends significantly along the trapping arm. A gradual increase in the equilibrium position is observed up to $x \approx -165\,\mu\text{m}$, where the ion height experiences a sharp rise, reaching $h \approx 170\,\mu\text{m}$.

Figure 4.4 presents the secular frequencies as a function of the ion position along the arm of the junction; again, the two shaded yellow regions indicate the portions of the junction arm corresponding to the perpendicular RF electrodes (darker area) and the central opening of the junction (lighter area). As discussed in Subsection 4.2.1, the central region of the junction is dominated by a hexapole field component rather than a pure quadrupole. This higher-order contribution provides weaker confinement, causing the secular frequencies to drop progressively as the ion approaches the junction center.

The first observation concerns the two radial secular frequencies, ω_y and ω_z. These frequencies remain nearly identical up to $x \approx -350\,\mu\text{m}$, as expected for a linear trap configuration, with small deviations likely arising from numerical fluctuations. Beyond this region, however, their behavior diverges. The z-axis frequency, ω_z, exhibits an initial drop and then stabilizes around 0.9 MHz near the trap center. This trend is consistent with the increase in ion–surface distance observed in panel (c) of Figure 4.3, indicating a locally weaker confinement along the z direction. Conversely, the y-axis frequency, ω_y, initially stabilizes slightly above 2 MHz before rapidly decreasing to 0.47 MHz as the ion approaches the center of the junction.

For completeness, the axial secular frequency ω_x is also reported. As discussed ear-

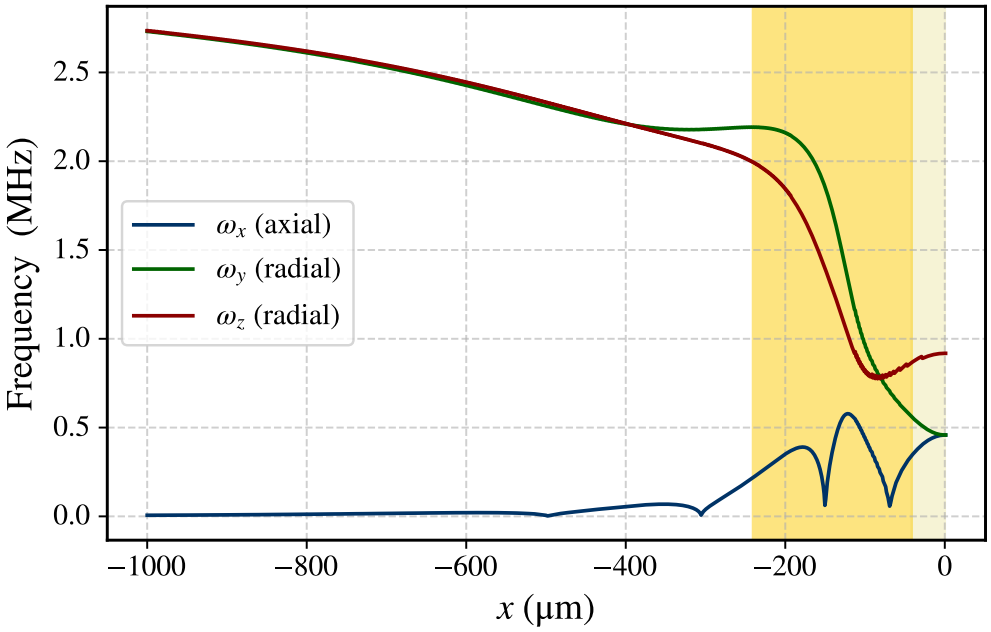

FIGURE 4.4: Secular frequency analysis of a non-optimized X-junction. The radial secular frequencies ω_y and ω_z remain nearly identical up to $x \approx -350\,\mu\text{m}$; then, ω_z decreases sharply and stabilizes at $\sim 1\,\text{MHz}$, while ω_y converges to the axial secular frequency ω_x, as a result of the weak axial confinement introduced by the perpendicular RF electrodes.

lier, the axial electric field is theoretically zero in the linear sections, implying that the corresponding confinement should vanish. Nevertheless, the results show a gradual increase of ω_x towards the junction center, where it reaches values comparable to ω_z. This behavior arises from the symmetry of the junction: if the simulations were performed along the perpendicular arm, the roles of ω_x and ω_y would be interchanged, and, for this reason, both converge to the same value at the trap center.

These results show that the non-optimized geometry suffers from several intrinsic limitations. Near the junction center, the secular frequencies decrease along all principal axes, indicating a loss of trap efficiency and weaker confinement. In addition, a pronounced pseudopotential barrier forms near the junction center, and the ion–surface distance varies significantly along the transport path. Together, these effects compromise the stability of the trapping potential and hinder reliable ion transport across the junction. Since the pseudopotential landscape depends solely on the geometry of the RF electrodes, these drawbacks can be addressed through an appropriate optimization of their shape. In the following section, we will focus on the junction optimization, showing how a refined electrode design can reduce the pseudopotential barrier, while maintaining a stronger radial confinement at the junction center.

4.3 X-junction optimization

The optimization of the junction geometry is a crucial step in improving the performance of ion transport, as its aim is to reduce the amplitude of RF electric fields and their gradients along the trap axis. Consequently, in an optimized junction design, the pseudopotential barrier and expected ion heating are reduced compared to the non-optimized version. In the following, we first describe the optimization procedure and its underlying theoretical principles. We then present the results obtained for the optimized geometries, followed by simulations of the corresponding heating rates.

4.3.1 Optimization principle and procedure

The optimization process aims at improving the junction performance by systematically adjusting the parameters that define its geometry. In general, this is achieved by introducing a *cost function*, a quantitative measure of how well a given geometry satisfies the desired criteria. The cost function includes all the quantities that should be minimized, such as the height of the pseudopotential barrier, variations in the ion–surface distance, or deviations from the target confinement strength. By varying the geometric parameters to minimize the cost function, the junction design progressively converges toward an optimal configuration.

The first step in optimizing the junction consists in establishing a parametric description of the electrode geometry. This is achieved by defining a set of control points along the borders of the RF electrodes, whose positions are varied during the optimization process [60, 79, 80, 84, 85]. The electrode edegs, and the overall junction geometry, can then be adjusted following two different approaches:

- *Piecewise parametrization*: the electrode edges are described by a series of straight segments connecting the control points. This approach offers high flexibility and simple geometric control, but can introduce discontinuities in the boundary curvature, leading to abrupt variations in the electric field.
- *Spline parametrization*: the electrode edges are defined by smooth, continuous spline curves passing through or near the control points. This allows for gradual field variations and smoother potential landscapes, at the cost of a more complex parameterization.

The piecewise approach presents two main drawbacks compared to the spline implementation. First, the presence of sharp corners at the control points can lead to localized field enhancement and an increased risk of RF breakdown, i.e. regions where the

electric field strength exceeds the breakdown threshold of the material or surrounding vacuum, potentially causing surface charging or discharging events that degrade the trap performance [80, 86, 87]. Second, the optimization must be constrained to prevent overlapping or intersecting RF electrodes. In the piecewise parametrization, ensuring this geometric consistency requires tighter restrictions on the allowed displacement of each control point, which in turn reduces the flexibility of the optimization and can slow down convergence towards an optimal solution. This is because small variations in the control point positions can induce abrupt changes in the pseudopotential landscape, making the cost function less smooth and therefore more difficult to minimize efficiently. The spline parametrization [44, 79, 80] overcomes these issues by imposing a continuous curvature of the electrode boundaries. This leads to a more regular pseudopotential landscape and a cost function that varies smoothly with the geometric parameters, resulting in improved numerical stability and more robust convergence during the optimization process, even if at the expense of higher computational cost. .

After defining the chosen parameterization, the optimization of the RF electrodes requires the simultaneous minimization of several objectives. A wide range of objective functions has been proposed in the literature [41, 60, 79, 80, 88], raising the question of which physical quantities are most relevant when designing an electrode structure for a specific application. In particular, previous studies have demonstrated that the presence of a moderate pseudopotential barrier does not necessarily preclude efficient, low-heating ion transport [41, 69]. For this reason, our optimization strategy aims to obtain a geometry that simultaneously balances multiple design requirements rather than solely eliminating the barrier.

To this end, we consider four optimization criteria, each quantifying a different aspect of the pseudopotential landscape relevant to stable and low-heating transport. These quantities are evaluated along the local minima of the pseudopotential:

$$r_{\min}(x) = \arg\min_{y,z} \Phi_{\mathrm{PS}}(x, y, z), \tag{4.4}$$

where arg min denotes the values of (y, z) that minimize the pseudopotential $\Phi_{\mathrm{PS}}(x, y, z)$ at a given axial position x. The optimization objectives we focused on are:

1. **Pseudopotential barrier**: The pseudopotential barrier, as previously discussed in Subsection 4.2.1, hinders ion transport through the junction. Minimizing it is therefore a primary objective in junction optimization [82], as it enables smoother ion shuttling and reduces motional excitation and micromotion. The correspond-

ing figure of merit is defined as:

$$F_1 = \int_{x_I}^{x_F} \Phi_{\mathrm{PS}}\big(r_{\min}(x)\big)\,\mathrm{d}x. \tag{4.5}$$

2. **Pseudopotential axial gradient:** The axial gradient of the pseudopotential can induce unwanted motional excitation of trapped ions. In particular, at junction intersections, where noise on the RF potential is non-negligible, these axial components may couple to the ions' motional modes, resulting in heating effects, as discussed in detail in Chapter 3. To suppress such effects, the term is added to the cost function:
$$F_2 = \int_{x_I}^{x_F} \left\| \frac{\partial}{\partial x} \Phi_{\mathrm{PS}}\big(r_{\min}(x)\big) \right\|^2 \mathrm{d}x. \tag{4.6}$$

3. **Variance of the confinement strength:** The radial potential curvatures determine the transverse confinement strength, as discussed in Subsection 2.1.4. Large variations in this curvature lead to changes in the secular frequencies, particularly near the junction center (see Figure 4.4 for a non-optimized geometry). When the confinement is weakened, the ions become more susceptible to drifts in stray electric fields [60] and to increased static motional heating caused by electric-field noise [44]. The considered mathematical form is [80]:
$$F_3 = \mathrm{Var}\left(\frac{\partial^2}{\partial y^2} \Phi_{\mathrm{PS}}\big(r_{\min}(x)\big) \right) + \mathrm{Var}\left(\frac{\partial^2}{\partial z^2} \Phi_{\mathrm{PS}}\big(r_{\min}(x)\big) \right). \tag{4.7}$$

4. **Ion–surface distance variations:** Variations in the trapping height must be carefully minimized to ensure uniform operation across the QCCD network. Significant deviations in the ion–surface distance can lead to ion losses at weakly confining intersections and inconsistencies during transport between connected zones of the network [79]. The corresponding figure of merit is defined as:
$$F_4 = \int_{x_I}^{x_F} \big[z_{\min}(x) - h\big]^2 \,\mathrm{d}x. \tag{4.8}$$

Therefore, the optimization can be formulated as a multi-objective problem, aiming to achieve a balanced trade-off among the various objective functions introduced above.

An efficient approach to solving multi-objective optimization problems is the *weighted sum method*, in which a scalar function U is minimized [89]:

$$U = \sum_{i=1}^{4} w_i\, F_i(d_1, d_2, \ldots, d_p) = \{w_1\xi_1, w_2\xi_2, w_3, \xi_3, w_4\xi_4\} \cdot \{F_1, F_2, F_3, F_4\} \tag{4.9}$$

Here, $w_i > 0$ is the weighting factor associated with the i-th objective function F_i, specifying its relative contribution according to the user's preference. To ensure that the objective functions F_i are of comparable magnitude in numerical calculations, the *lower-bound transformation method* [89] was applied. The corresponding transformation coefficient ξ_i is defined as:

$$\xi_i = \frac{1}{F_i^0}\,, \tag{4.10}$$

with:

$$F_i^0 = \min\{F_i(d_1, d_2, \ldots, d_p)\}\,, \tag{4.11}$$

where $d_1, \ldots, d_p$ denote the p degrees of freedom of the chosen parametrization. Notice that we normalize using the minimum value rather than the maximum because the maximum of F_i could be far from the optimal region and would artificially distort the relative weighting of the objectives. Additionally, it is important to distinguish between the roles of the weights w_i and the transformation coefficients ξ_i. The weights w_i encode the user's preferences and are chosen as positive values within the interval $(0, 1)$ to specify the relative importance of each objective. In contrast, the transformation coefficients ξ_i are used to scale the objective functions so that they have comparable magnitudes in the numerical calculations [79].

Before presenting the optimization results, it is important to note that all simulations were carried out using *normalized* (non-dimensional) objective functions F_i. In particular, the pseudopotential was expressed in non-dimensional form as:

$$\tilde{\Phi}_{\mathrm{ps}} = \Phi_{\mathrm{PS}} \left(\frac{Q^2 V_{\mathrm{RF}}^2}{4m\,\omega_{\mathrm{RF}}^2\, h^2} \right)^{-1} = \left\| \nabla \Theta_{\mathrm{RF}} \right\|^2 , \tag{4.12}$$

where Θ_{RF} denotes the RF-electrode basis function [90], which depends solely on the trap geometry. Specifically, the basis function for the i-th electrode is defined as:

$$\Theta_i(\mathbf{x}) = -\frac{1}{4\pi} \oint_{S_i} \frac{\partial G(\mathbf{x}, \mathbf{x}')}{\partial n'}\, da'\,, \tag{4.13}$$

where $G(\mathbf{x}, \mathbf{x}')$ is the Green's function for Laplace's equation (Equation 2.2), S_i denotes the surface of the i-th electrode, and $\partial/\partial n'$ indicates the derivative along the outward normal at $\mathbf{x}'$. The function Θ_i represents the electric potential produced when the i-th electrode is held at 1 V while all other electrodes are grounded [90].

Within this framework, both the objective functions and the resulting optimized solutions become independent of the ion mass, the trapping parameters (V_{RF} and ω_{RF}), and

the trapping height. As a result, the optimization focuses purely on the geometric properties of the trap, independent of any specific physical scale. Once the optimization is completed, the pseudopotential and the other quantities are transformed back, in order to recover physical quantities. Consequently, the optimized geometries obtained in this work can be regarded as *universal* surface-electrode configurations, suitable for any ion species or set of trapping parameters.

The following subsections present and discuss in detail the results obtained for the two electrode parametrizations introduced earlier. All simulations are performed on an X-junction with overall dimensions of 2500 µm × 2500 µm and an electrode separation of 5 µm, which allows us to work in the gapless-approximation presented in Subsection 2.2.2. The optimization region, covering 1500 µm × 1500 µm, is chosen to ensure that boundary effects near the outer trap edges do not affect the results, thereby preserving the reliability of the optimization within the region of interest.

4.3.2 Piecewise parametrization

We begin our analysis by introducing the piecewise parametrization of the junction electrodes. In this approach, the geometry of each electrode is defined by a discrete set of control points, whose coordinates serve as inputs to the multi-objective procedure described in Subsection 4.3.1 and governed by Equation 4.9. The number and placement of these points must be carefully chosen: too few points limit the flexibility of the electrode geometry, potentially resulting in suboptimal junction designs, while too many points can introduce numerical instabilities and significantly increase the computational cost. It should be noted that different control-point configurations can result in distinct optimized geometries, and the choice of the weights w_i in the multi-objective function, which reflect the relative importance assigned to each objective, can further influence the resulting electrode shapes. Here, we set $w_i = 1 \ \forall i$, allowing the optimization to pursue a balanced trade-off among the objectives considered. For clarity, we present only a single representative example of the resulting optimized geometry.

Owing to the rotational and reflection symmetry of the junction, both about the central axis and under 90° rotations, it is sufficient to perform the optimization on a single arm of a single quadrant of the trap, simulating ion shuttling along the transport channel defined in Equation 4.4. Each L-shaped RF electrode has two independent edges, an inner and an outer one, both of which are subject to optimization. In the present work, the control points are distributed equidistantly along the edges of the RF electrodes. This distribution provides sufficient flexibility to optimize the electrode shape while avoiding

arbitrary assumptions about which regions of the junction are most critical. For each edge, seven control points with fixed x-coordinates and variable y-coordinates were employed, allowing for trap designs that respect geometric and physical constraints, preventing unphysical configurations such as self-intersecting polygons. In addition, one extra control point per edge was placed along the bisector of the quadrant to ensure smoothness near the junction center and preserve the symmetry.

The initial and optimized trap geometries are compared in Figure 4.5. Panel (a) shows the initial, non-optimized junction together with the set of control points defining the electrode boundaries. Panel (b) displays the optimized geometry obtained from the piecewise parametrization. The inner edge exhibits minimal variation, except for the control point corresponding to the vertex, which shifts closer to the junction center. The outer border, in contrast, undergoes a more substantial deformation, resulting in a peculiar shape that best satisfies the multi-objective optimization criteria.

The different trap geometries result in different pseudopotential landscapes. To quantify this effect, Figure 4.6 reports the pseudopotential, its gradient, and the ion–surface distance along the left arm of the junction. For comparison, the results for the non-optimized junction are also shown as dashed lines, providing a reference for the impact of the optimization. Panel (a) shows that the pseudopotential barrier is reduced by approximately one third ($\sim$ 3.7 meV), and the maximum is shifted closer to the junction

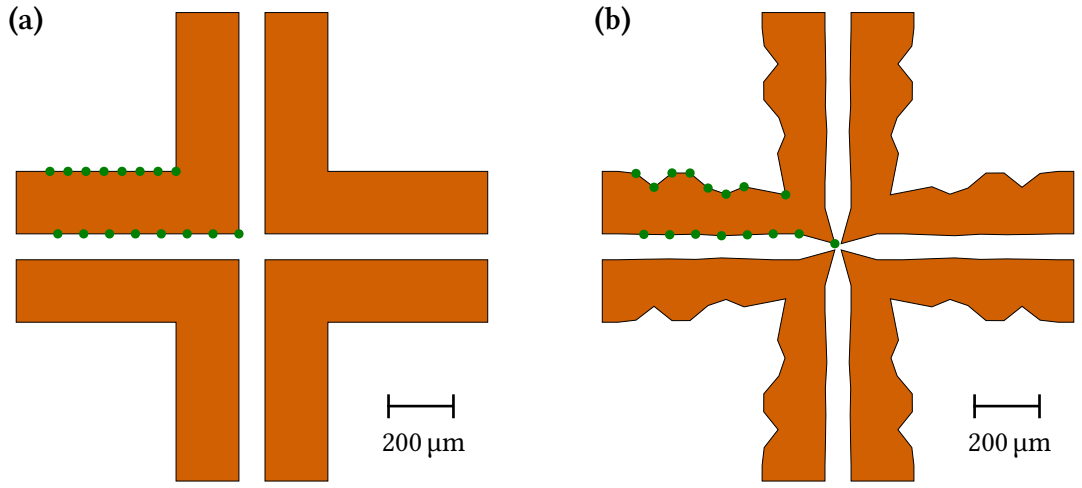

FIGURE 4.5: Comparison between non-optimized and piecewise-optimized junction geometries. Panel (a) shows the initial configuration along with the control points; panel (b) displays the optimized geometry, where the inner border remains nearly unchanged while the outer border deforms significantly to satisfy the multi-objective criteria.

center. This is advantageous, as it mitigates the influence of the junction on the linear segments of the trap. Panel (b) displays the gradient, which does not change significantly in amplitude but is shifted toward the junction center, as expected from the results shown in Panel (a). As a result, the region in which the ion is more subjected to heating is smaller compared to the unoptimized version of the trap. Finally, panel (c) illustrates the ion–surface distance as a function of the position along the junction. The ion height remains relatively constant during shuttling, maintaining a value close to the target trap height compared to the non-optimized case. However, when approaching the junction center ($-100\,\mu\text{m} \leq x \leq -35\,\mu\text{m}$), the ion must be lowered to remain at the minimum of the pseudopotential, reaching a minimum ion height of $h_{\text{min}} = 89\,\mu\text{m}$. As

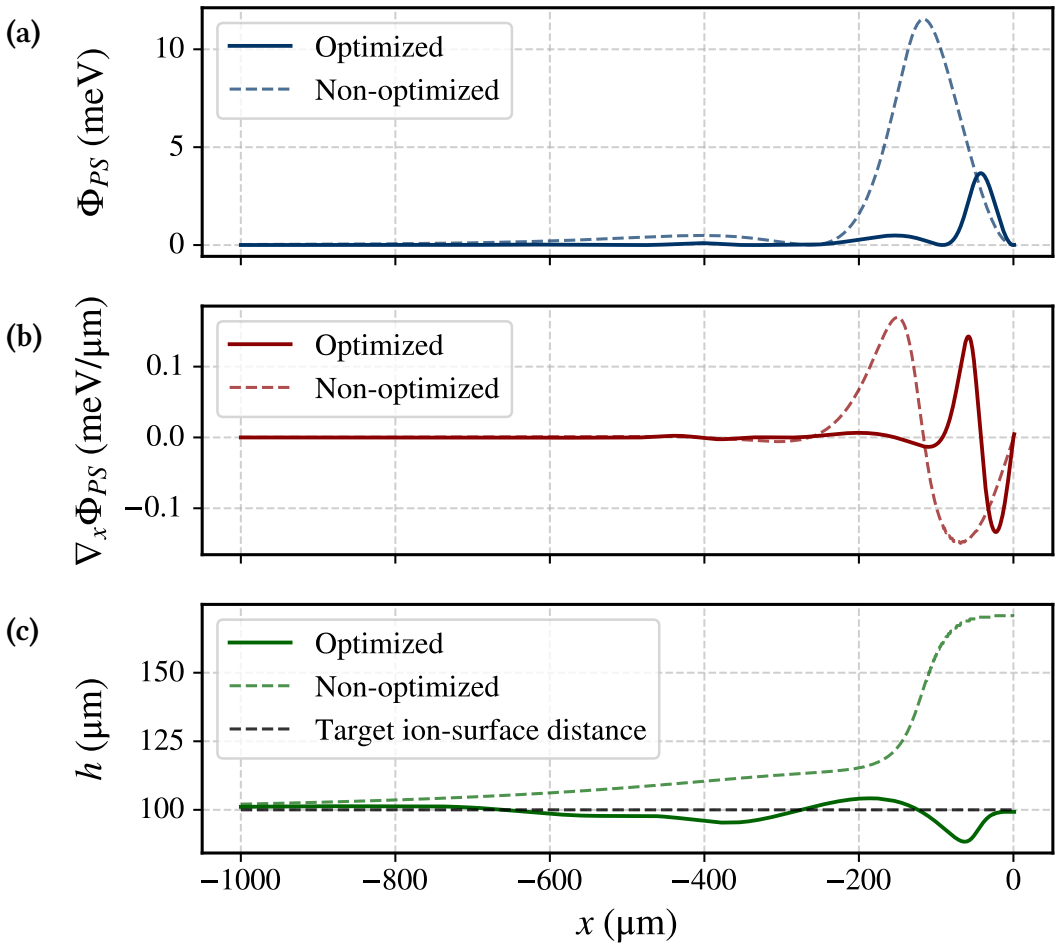

FIGURE 4.6: Pseudopotential analysis along one arm of the piecewise parametrized optimized X-junction. (a) Pseudopotential barrier is reduced and shifted toward the junction center. (b) Gradient is slightly shifted with similar amplitude, reducing the region prone to heating. (c) Ion height remains near the target along the junction, with minor adjustments near the center to follow the pseudopotential minimum.

the ion approaches the junction center, its height is increased back to h to ensure that it remains aligned with the pseudopotential minimum.

The other key quantity that characterizes the trap performance is the confinement itself, which is directly captured by the secular frequencies. The comparison between the optimized and non-optimized case is shown in Figure 4.7. The optimization clearly enhances confinement: both radial frequencies remain stable along most of the transport path, with a sudden drop around $x \approx -95\,\mu\text{m}$. Despite this drop, ω_z quickly stabilizes near 1.5 MHz (compared to ~ 0.9 MHz in the non-optimized case), ensuring tighter confinement along the z-axis. Meanwhile, ω_y reaches ω_x at roughly 0.75 MHz. The axial frequency produced by the perpendicular RF electrode only becomes significant very close to the junction center ($x \lesssim -100\,\mu\text{m}$), highlighting that the optimized junction maintains strong radial confinement along the transport path while postponing axial effects until the ion arrives near the junction center.

Overall, the piecewise parametrization demonstrates a clear improvement over the non-optimized design, providing a more controlled and efficient ion transport. By adjusting the electrode geometry, the optimization reshapes the X-junction pseudopotential landscape: the barrier is lowered and shifted toward the junction center, the pseudopotential gradient is repositioned so that the regions where ions are more prone to

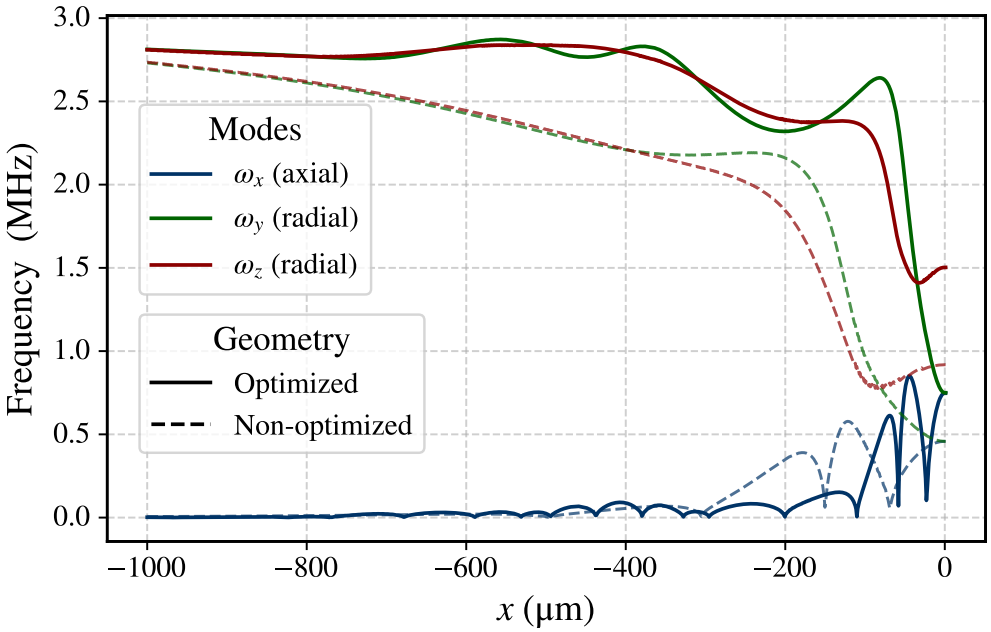

FIGURE 4.7: Comparison of secular frequencies between non-optimized and piecewise-optimized junction geometries. The optimized junction shows improved radial confinement along most of the transport path, while axial effects become significant only near the junction center.

heating are reduced, and the ion-surface distance remains more stable along the transport path, with variations limited to within 10% of the target value. These improvements are reflected in the secular frequencies as well, with radial confinement maintained along most of the junction, while the axial frequency varies from 0 kHz to 800 kHz as the ion is moved closer to the trap center. In the next section, we turn to the spline parametrization, exploring an alternative strategy for optimizing the junction geometry.

4.3.3 Spline parametrization

The spline parametrization of the junction is implemented by describing the inner and outer boundaries of the RF electrodes using cubic B-spline functions [79, 80]. As before, each edge is defined by eight control points and exhibits reflection symmetry with respect to the diagonal axis of the junction. During the numerical optimization, each spline is sampled at one hundred evenly spaced points along its parameter domain, and the resulting points are connected by straight line segments to form the electrode polygons. As in the previous case, all objective weights are set to $w_i = 1 \ \forall i$, ensuring a balanced optimization across the selected criteria. The optimization procedure follows the same principles outlined earlier, simulating ion shuttling along the transport channel defined in Equation 4.4, while allowing the y-coordinates of the control points to vary within geometric constraints that prevent unphysical configurations such as self-intersecting polygons.

A representative result of the optimized junction geometry obtained with the spline parametrization is shown in Figure 4.8. Panel (a) illustrates the initial, non-optimized junction together with the control points defining the electrode edges. The same set of control points used in the piecewise parametrization is adopted here as the starting configuration for the optimization. Panel (b) presents the optimized configuration, where the inner edge remains nearly unchanged in the y-direction, except for the vertex control point, which shifts closer to the junction center. The outer edge also undergoes a smooth yet more intricate deformation, particularly noticeable near the diagonal. Overall, the resulting contour closely resembles the one obtained with the piecewise parametrization shown in Figure 4.5; the main difference lies in the nature of the spline representation, which produces a much smoother profile while preserving the same underlying geometric trend.

The geometric refinements introduced by the spline parametrization have a clear impact on the pseudopotential landscape along the transport path. Figure 4.9 illustrates

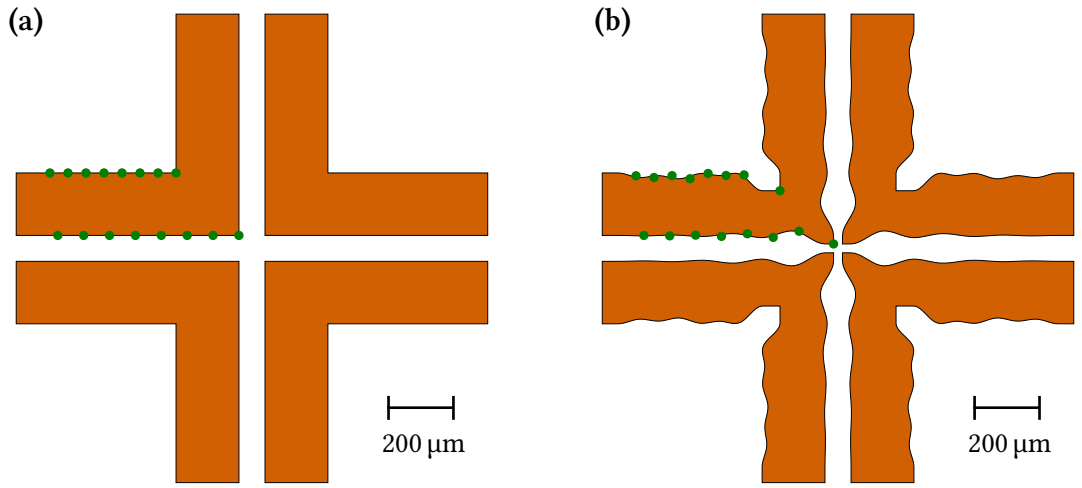

FIGURE 4.8: Comparison between non-optimized and spline-optimized junction geometries. Panel (a) shows the initial configuration along with the control points; panel (b) displays the optimized geometry.

these effects by showing the pseudopotential, its gradient, and the ion–surface distance along the left arm of the junction. Dashed lines indicate the non-optimized junction, serving as a reference. Panel (a) shows that the pseudopotential exhibits small bumps beginning around $x \approx -500\,\mu\text{m}$, reaching amplitudes of approximately 0.5 meV. The central barrier is reduced to about one tenth of that in the non-optimized case, with a maximum value of $\sim$ 1.8 meV, and is shifted closer to the junction center, indicating a clear improvement over the piecewise parametrization. Panel (b) depicts the gradient: although it exhibits more oscillations than in both the non-optimized and piecewise cases, its overall amplitude is smaller. Consequently, for the same noise amplitude, the ion heating would be reduced. Nevertheless, the pseudopotential gradient is different from zero in a region which is bigger compared to the case of the piecewise parametrization, extending the spatial extent over which heating effects may occurr. Panel (c) presents the ion–surface distance as a function of the position along the junction, which fluctuates more than in the piecewise parametrization. In particular, it follows the non-optimized trajectory up to $x \approx -150\,\mu\text{m}$, after which it decreases to remain aligned with the pseudopotential minimum, reaching a minimum value of approximately 91 µm around $x \approx -50\,\mu\text{m}$. It then rises again, returning to about 97 µm at the junction center. Overall, the spline parametrization produces a lower pseudopotential barrier and a reduced gradient compared to both the piecewise parametrization and

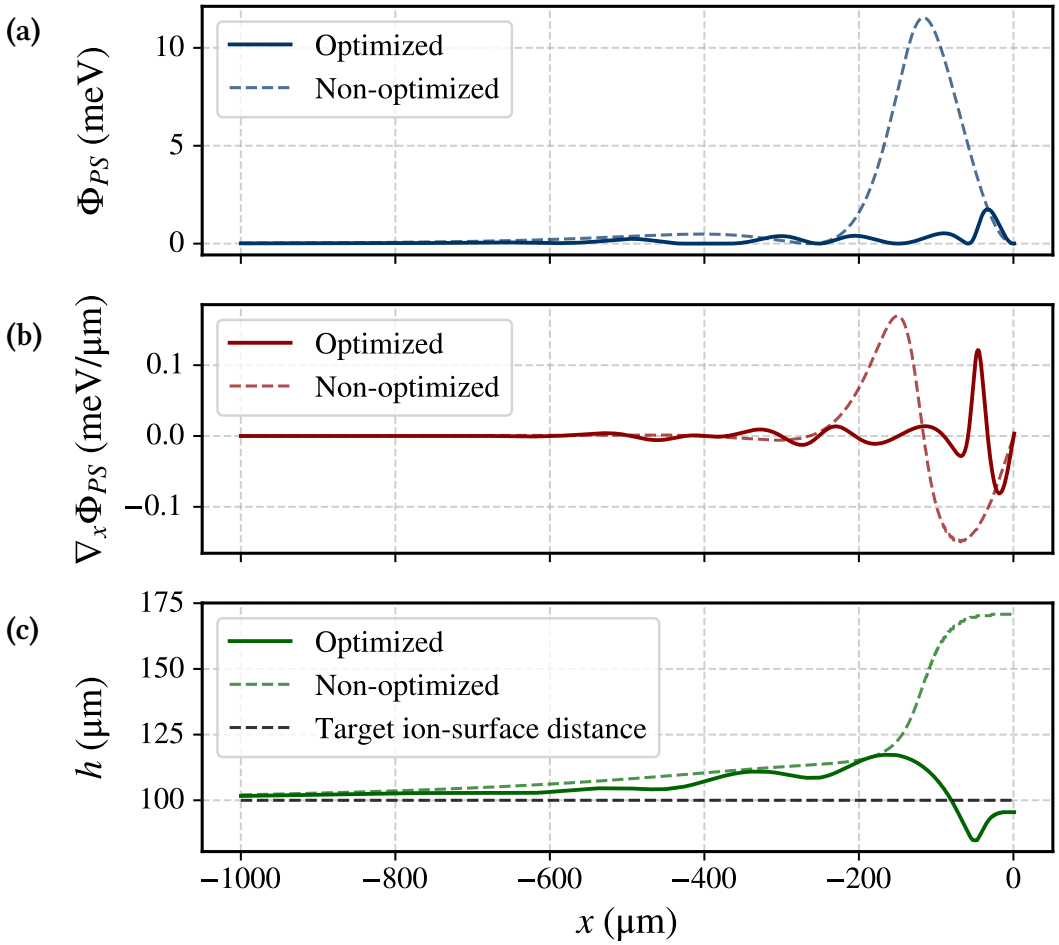

Figure 4.9: Pseudopotential analysis along one arm of the spline parametrized optimized X-junction. The pseudopotential (a), its gradient (b), and the ion–surface distance (c) are shown, with dashed lines indicating the non-optimized junction. Compared to the piecewise parametrization (Figure 4.6), the spline approach reduces the pseudopotential barrier and the gradient, ensuring smoother transport and lower heating, even though the ion–surface distance exhibits larger variations.

the non optimized geometries, ensuring smooth ion transport while mitigating heating from technical noise.

Concerning the secular frequencies, Figure 4.10 compares the spline-optimized junction to the non-optimized case. Unlike the piecewise parametrization (Figure 4.7), the spline does not significantly enhance confinement: the radial frequencies largely follow the non-optimized trend, with a brief drop near $x \approx -50\,\mu\text{m}$. Afterwards, ω_z rises to about 1.3 MHz (compared to $\sim$ 0.9 MHz for the non-optimized and $\sim$ 1.5 MHz for the piecewise case), while ω_y approaches ω_x around 0.75 MHz, similarly to the piecewise junction. The axial contribution from the perpendicular RF electrode begins farther

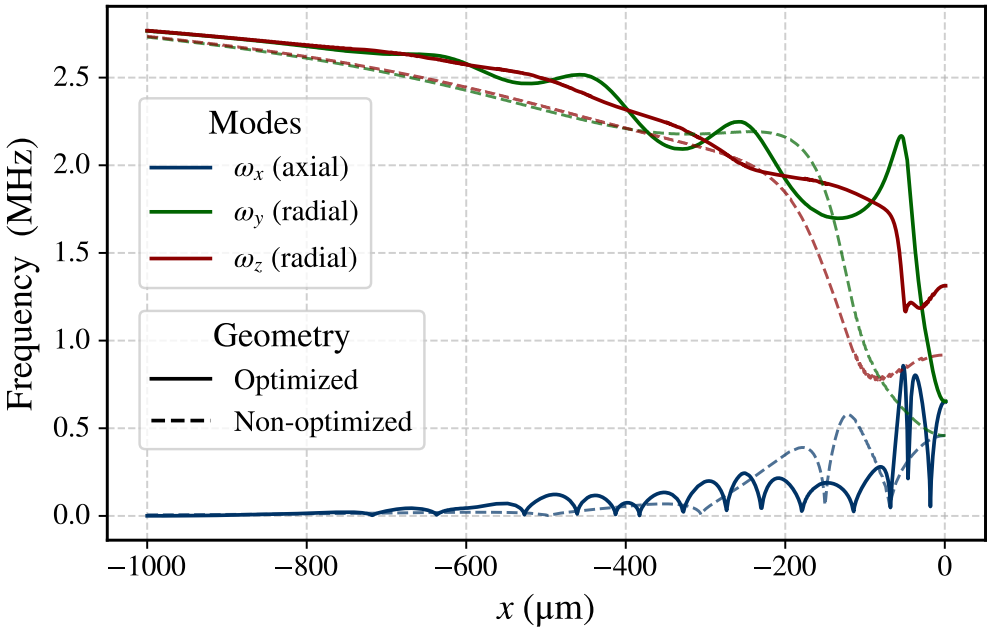

Figure 4.10: Comparison of secular frequencies for spline optimized and non-optimized junctions. Radial confinement largely follows the non-optimized case, with a brief drop near $x \approx -50\,\mu\text{m}$, while ω_z rises to 1.3 MHz. Compared to the piecewise parametrization (Figure 4.7), the spline maintains radial confinement and delays axial effects, but achieves slightly weaker overall confinement.

from the junction center but only becomes significant close to $x \approx -100\,\mu\text{m}$. Overall, the spline optimization preserves radial confinement along most of the transport path and delays unwanted effects along the axial direction, even though the piecewise parametrization achieves stronger confinement along the z-axis.

From these observations, we conclude that the spline parametrization significantly improves the non-optimized design by lowering the pseudopotential barrier and reducing the gradient, which together enable smoother ion transport and mitigate heating from technical noise. While radial confinement remains largely preserved along the transport path, the ion–surface distance fluctuates more than in the piecewise case, and the region affected by gradient oscillations is extended by 50%. Undesired effects along the axial direction begin to appear from about $x \approx -120\,\mu\text{m}$ as the ion approaches the junction center, although the overall z-axis confinement remains moderatly less tight than in the piecewise parametrization. Overall, the spline optimization offers an alternative approach to junction shaping, effectively balancing a reduced pseudopotential barrier and heating with moderately slightly weaker confinement.

4.4 Ion motional heating in the X-junction

Having characterized the X-junction and the optimization strategies, we now turn our attention to ion heating. The optimized junction geometries are expected to reduce heating by producing a smoother pseudopotential landscape and more uniform secular frequencies, as seen in Section 4.3, but a quantitative assessment is necessary to confirm this gain. The aim of this section is to apply the noise models developed in Chapter 3 to the junction geometries previously presented, in order to estimate the corresponding heating rates. This analysis provides a comprehensive characterization of ion heating in the X-junction for the QCCD architecture, illustrating how different electrode parametrizations can affect the stability and robustness of ion transport.

The simulations and the optimization of the junction, as recalled at the end of Subsection 4.2.1, are carried out with a fixed RF voltage of $V_{\mathrm{RF}} = 42.2\,\mathrm{V}$. As the junction is driven with a single RF signal, phase noise does not contribute to ion heating, as discussed in Subsection 3.2.2. Thus, for X-junctions, RF amplitude noise is the only mechanism related to the RF drive capable of inducing ion heating. In this case, the expression for the heating rate takes the simple form given by Equation 3.20, repeated here for clarity:

$$\dot{\bar{n}} = \frac{Q^4}{16m^3\Omega_{\mathrm{RF}}^4\hbar\omega}\left[\vec{\nabla}\left(\left|\vec{E}_0\right|^2\right)\right]^2 \frac{S_{V_N}(\Omega_{\mathrm{RF}} \pm \omega)}{V_{\mathrm{RF}}^2}, \tag{4.14}$$

where $S_{V_N}(\Omega_{\mathrm{RF}} \pm \omega)$ is the spectral voltage noise density. As the value of $S_{V_N}(\Omega_{\mathrm{RF}} \pm \omega)$ is dependent on the experimental apparatus used to drive the trap, one could either assume a plausible value or express the results in a normalized form as $\dot{\bar{n}}/S_{V_N}(\Omega_{\mathrm{RF}} \pm \omega)$, the latter being the approach adopted in this work.

To obtain physically meaningful heating rates, the ion must remain confined during shuttling along the junction. Thus, we will consider the axial secular frequency to be fixed to 1 MHz along the junction, a value known to provide sufficient confinement [41, 61, 80], while the radial frequencies follow the behaviour resulting from the simulations. A constant axial trapping frequency can be obtained adjusting the axial field curvature with the DC electrodes along the shuttling path. The analysis of the required DC voltage waveforms for all junction configurations is presented in Section A.2.

Heating rates in the non-optimized X-junction

Figure 4.11 presents the results of the heating rate simulations performed for the non-optimized junction. The shuttling trajectory was restricted to the interval $-500\,\mu\mathrm{m} \leq$

$x \leq 0\,\mu\text{m}$, since significant variations in the pseudopotential gradient are observed only in the vicinity of the junction center (see Figure 4.3). For comparison, the pseudopotential profile is also reported in the same figure, overlaid in the background.

While all three components of $\dot{\bar{n}}$ are plotted, along the radial directions the ion barely heats up compared to the axial direction. This behaviour is consistent with expectations. The symmetry of the system with respect to the y-axis makes the gradient of the pseudopotential along this direction identically zero, whereas the z-component, although not strictly null, remains very low due to the small electric field gradient in that direction and the limited displacement of the ion along z. Thus, only the heating along the x-direction is of interest.

The axial heating rate exhibits two distinct peaks, which are located at approximately $x \approx -150\,\mu\text{m}$ and $x \approx -60\,\mu\text{m}$, corresponding to the rise and fall of the pseudopotential barrier generated by the perpendicular RF electrode (see again Figure 4.3). The second peak reaches a slightly lower maximum value than the first, reflecting the small asymmetry of the barrier shape. It is worth noting that, since the heating rate depends on the spatial gradient of the pseudopotential, minimal heating occurs at the top of the barrier, despite this being the point of maximum axial RF field and ion micromotion. As the ion

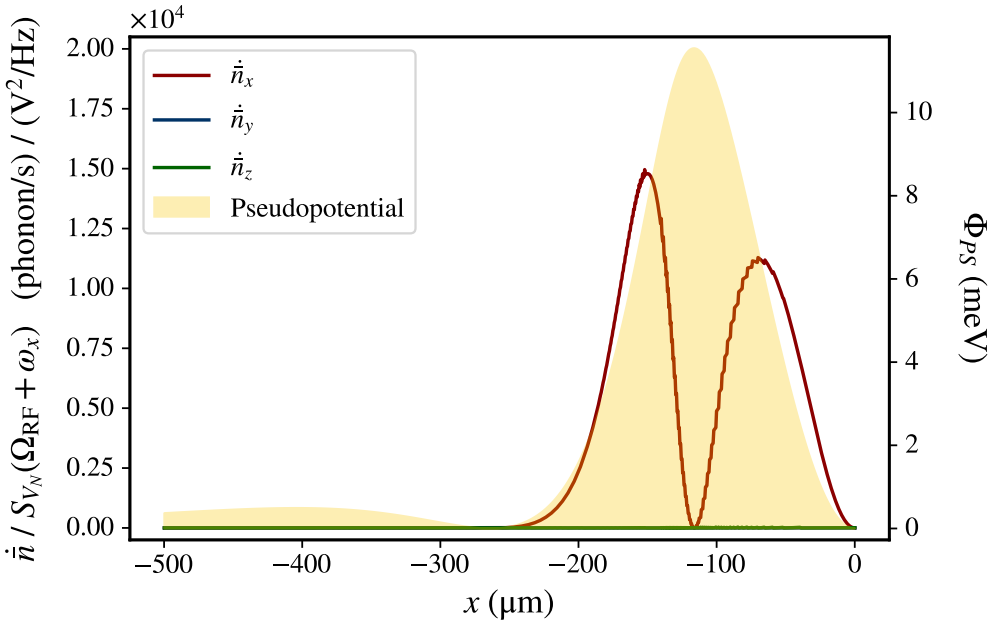

Figure 4.11: Heating rate simulation for the non-optimized junction. The pseudopotential profile is shown in the background for reference. Only the axial component of the heating rate exhibits significant values, with two peaks around $x \approx -150\,\mu\text{m}$ and $x \approx -60\,\mu\text{m}$ corresponding to the rise and fall of the pseudopotential barrier near the perpendicular RF electrode.

approaches the junction center, the heating rate gradually decreases to zero.

In summary, each arm of the non-optimized junction presents two critical points that can lead to unwanted motional heating during ion transport, which correspond to the transitions between the linear region and the junction center, where the pseudopotential varies most rapidly.

Heating rates in the optimized X-junction

The expression for the heating rate given in Equation 4.14 remains valid for the optimized junction. In both the piecewise and spline parametrizations, a reduction in the pseudopotential gradient is observed, together with a shift of its maximum closer to the junction center (see Figure 4.6 and Figure 4.9, respectively). Consequently, for the same noise conditions, a corresponding reduction and spatial shift in the heating rate are expected. Figure 4.12 reports the simulated heating rates for both parametrizations, with the corresponding results for the non-optimized junction included for comparison.

Panel (a) shows the results for the piecewise parametrization. Two peaks are again observed, located approximately at $x \approx -60\,\mu\text{m}$ and $x \approx -20\,\mu\text{m}$, both shifted towards the junction center compared to the non-optimized case. Their maximum amplitudes are reduced, reaching about 72% of the value observed before and they exhibit a smaller difference in amplitude than in the case of the non-optimized junction.

Panel (b) shows the results for the spline parametrization. In this case, the two peaks are shifted even further towards the junction center, appearing at $x \approx -50\,\mu\text{m}$ and $x \approx -15\,\mu\text{m}$. Both peaks exhibit a lower amplitude than in the previous cases, with the maximum reaching roughly half of that in the non-optimized junction. The asymmetry between the two peaks becomes more pronounced here, with the second peak reaching only about 45% of the amplitude of the first.

Both parametrizations result in a significant reduction of the heating rate, indicating a lower susceptibility of the ion to RF technical noise during shuttling. Thus, ion transport through the optimized junctions is expected to exhibit higher stability and fidelity, with reduced motional excitation and improved overall performance of the shuttling process.

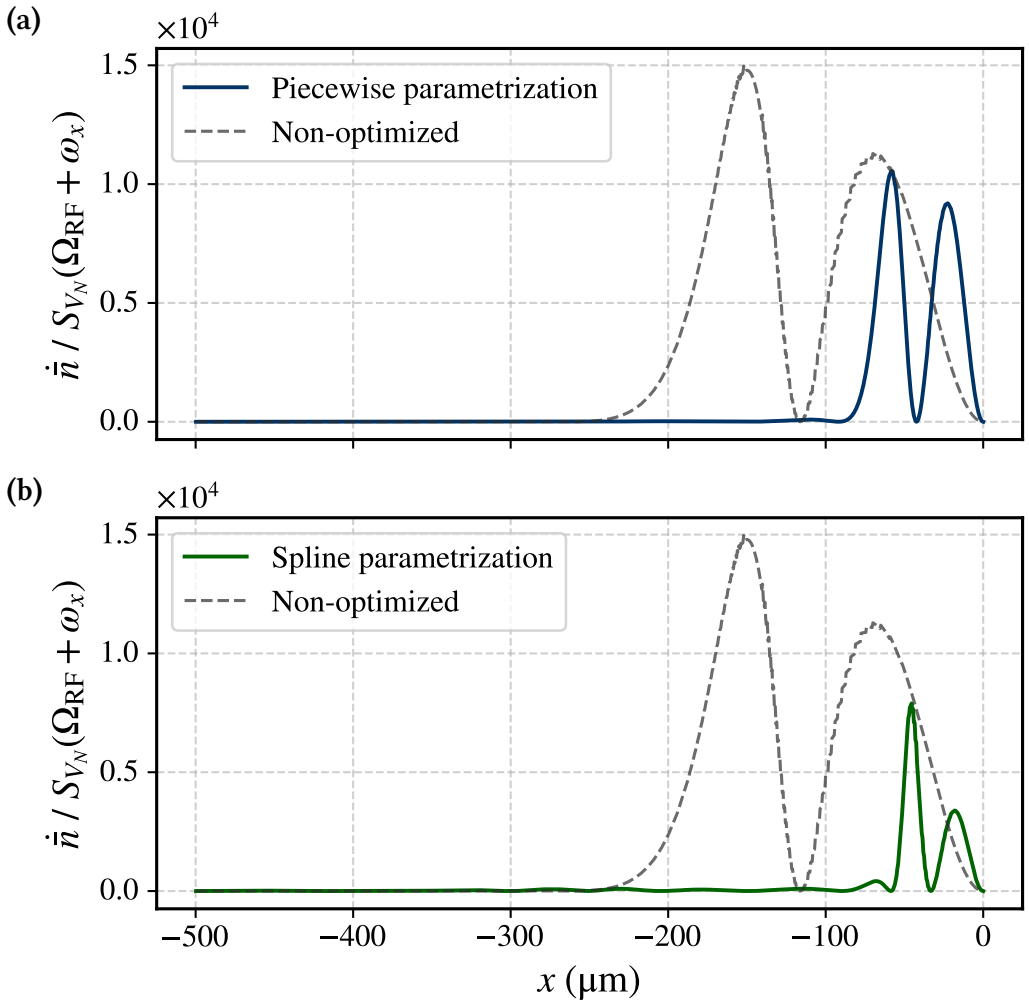

FIGURE 4.12: Simulated heating rates for the optimized junction using the piecewise (a) and spline (b) parametrizations. In both cases, two peaks are observed, shifted towards the junction center compared to the non-optimized junction. For the piecewise case, both peaks exhibit reduced and nearly symmetric amplitudes, while in the spline case they are further shifted and significantly weaker, with a more pronounced asymmetry between them.

5

RF technical noise in the QSA architecture

In this chapter, we investigate the effects of RF technical noise during RF transport in the QSA architecture. The first section introduces the fundamental concepts of the QSA architecture. The second section provides a characterization of the trap used to study RF transport, serving as a testbed for its later implementation in the broader QSA framework. After a brief description, the pseudopotential analysis is presented. The third section introduces heating rate simulations obtained by applying the noise models developed in Chapter 3 to the radial ion transport. Finally, the last section compares the simulations with the experimental data collected by Marco Valentini and Martin van Mourik. After a brief description of the experimental setup, the ion heating rate measurements are presented and compared with the developed models.

5.1 The QSA architecture

The QSA architecture [45] represents a promising alternative to the QCCD approach. Instead of a collection of ions distributed across separate zones and transported along complex paths, the QSA arranges ions into a two-dimensional lattice of coupled ion strings. This configuration can be viewed as an array of parallel linear traps, each hosting a small ion chain with its own trapping parameters. The central idea is that every lattice site contains a short ion string acting as a local quantum register, while neighbouring sites are electrostatically coupled through precisely engineered potentials that control the interaction strength between chains.

Figure 5.1 schematically illustrates the working principles underlying the QSA architecture. Panel (a) shows a single trapping site, or qubit register, consisting of a few ions

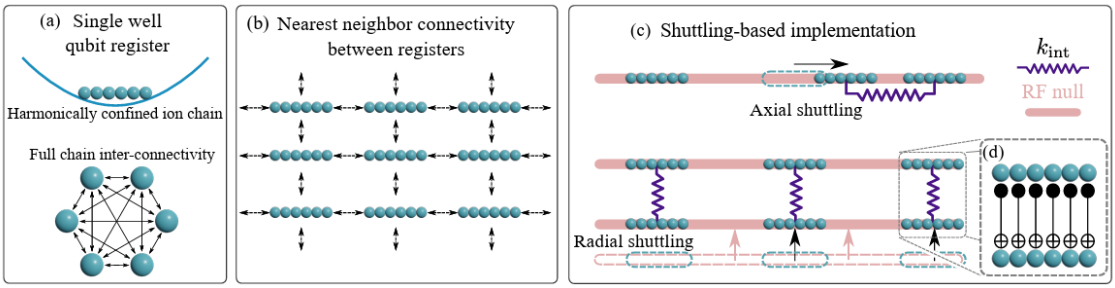

FIGURE 5.1: Overview of the main features of the proposed QSA architecture. (a) A chain of ionic qubits are confined in a single well of an ion trap, with full connectivity between all qubits. (b) Multiple separate trapping regions are distributed over a 2D lattice. Trapping sites have connectivity between nearest neighbors. (c) A shuttling-based approach is used to bring ion chains close enough to couple them, with an interaction constant k_{int}, symbolically depicted in purple as springs. Axial shuttling is achieved by displacing an ion string along a linear trap's RF null, and radial shuttling is achieved by displacing the RF null itself. Image courtesy of Marco Valentini, adopted from [45].

harmonically confined in a single potential well. Every ion within the string can be individually addressed and entangled with any other ion through the shared motional modes of the chain. The full connectivity available within each trapping site enables the implementation of high-fidelity local gate operations and the realization of error detection and correction routines before interacting with neighboring chains. Panels (b) and (c) highlight QSA's defining feature: conversely to the QCCD architecture, which is based on splitting and merging operations of ion registers, in the QSA architecture, information is processed by keeping ions in separate wells. Distinct ion registers can interact via the Coulomb force, which scales with the inter-well distance. The QSA thus relies on the tuning of the distance between ion registers to either enhance or suppress the coupling between them. When two neighboring registers are brought sufficiently close, their motional modes hybridize, giving rise to an effective interaction of strength k_{int}, depicted symbolically as springs in panel (c), hence the term *spring array*. This configuration enables all-to-all connectivity between ions belonging to adjacent strings, enabling the realization of parallel, transversal gate operations across multiple registers, as illustrated in panel (d).

Two dimensional connectivity across the entire lattice requires the capability to tune the position of each well in two dimensions. Along the trap axis, this is achieved via DC transport operations. Along the direction perpendicular to it, one has to shift the position of the RF null itself. This requires a different kind of transport strategy, which is referred to as *RF transport*, as it is realized by shaping the pseudopotential landscape controlling the voltage applied to the RF electrodes.

From an experimental perspective, the QSA architecture has demonstrated the controlled formation of two-dimensional ion lattices and coherent motional coupling between neighboring strings [45]. However, the realization of coherent operations in combination with RF transport has not yet been achieved, primarily due to the high ion heating rates observed during radial shuttling. In this chapter, we apply the theoretical models developed in the previous chapters to a specific trap design, used as a testbed for RF transport. Through this analysis, we aim to identify the physical mechanisms and design parameters that determine the heating rate, clarifying which aspects of the system could be adjusted to enable low-heating, high-fidelity RF shuttling.

5.2 RF transport: trap characterization

As previosuly discussed, in the QSA architecture, ion registers are arranged in a 2D lattice, where the connectivity across the lattice is achieved via ion transport. Along the axial direction of the trap, ion transport is realized controlling the voltage applied to the DC electrodes, while along the radial direction, the position of ion registers is tuned controlling the RF voltage applied to the trap. DC shuttling is a well-established technique, extensively discussed in the literature [20, 59, 91, 92], and will therefore not be treated further in this work. In contrast, RF pseudopotential control has only been explored in two prior experiments, where instabilities in the RF electronics and limited DC control prevented the successful coupling of ions located in separate linear traps [93, 94]. To address these challenges, a dedicated device, the 4S trap (abbreviation for *Quattro Stagioni* trap) was specifically designed to demonstrate controllable radial transport and programmable well-to-well coupling along the radial direction.

5.2.1 Trap design

The 4S trap, shown in Figure 5.2, features three parallel RF electrodes on the surface, referred to as the inner (RF1) and outer (RF2) electrodes, with respective widths of 75 µm and 255 µm, separated by 115 µm. The geometry of these RF electrodes is derived from the BB trap originally developed by Philip Holz [95]. For clarity, throughout this discussion we will refer to the coordinate system shown in Figure 5.2. In this reference frame, the z-axis corresponds to the axial direction of the trap (along the ion chain), while the x- and y-axes define the radial plane. Specifically, y denotes the vertical direction (height above the electrode surface) and x the horizontal direction (width across the surface).

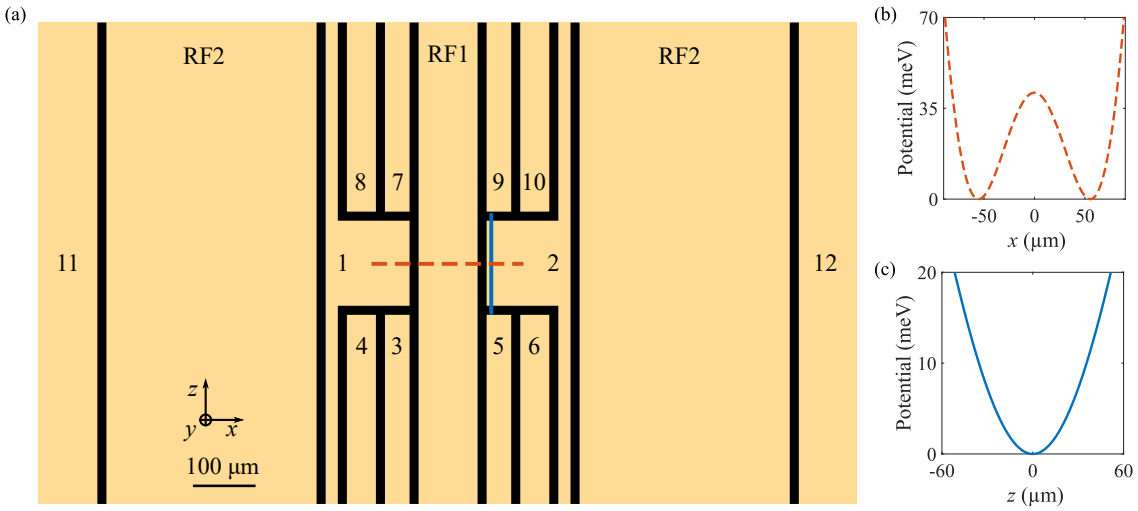

FIGURE 5.2: 4S trap electrode layout and electric potentials. (a) Schematic of the trap center. The inner and outer RF electrodes are labelled RF1 and RF2 respectively, as they can be at different voltages. The electrodes in gray are the DC electrodes used for confinement and control of the two trapping sites. Two additional DC electrodes, not shown in the image, are present alongside the outer RF electrodes. (b) Cross section of the total potential along the radial direction x corresponding to the orange dashed line in (a) generated by applying a RF voltage of 215 V peak-to-peak at 19 MHz on the RF electrodes. The resulting radial frequency along x is 2.4 MHz in both trapping sites. (c) Cross section of the total potential along the axial direction z (blue line in (a)), with an axial frequency of 1 MHz. Such axial frequency is generated with a voltage set with a maximum DC amplitude of 1.3 V. Image courtesy of Marco Valentini.

As discussed in Section 5.1, ion strings are confined at the minimum of the effective trapping potential generated by the RF electric field, commonly referred to as the null of the RF pseudopotential. In the configuration shown in panel (a) of Figure 5.2, the presence of three coplanar RF electrodes gives rise to a pseudopotential landscape characterized by two distinct minima, as illustrated in panel (b), which can be independently controlled via 12 DC electrodes, as shown in panel (c). The 4S trap was fabricated by Infineon Technologies Austria through the deposition of a single 2 µm aluminum layer onto a 725 µm-thick, double-sided polished fused silica substrate.

5.2.2 RF PSEUDOPOTENTIAL CONTROL

The 4S trap is designed to offer independent control over two trapping sites located in distinct linear traps, while also allowing their separation to be tuned via precise adjustment of the RF voltages. Because two-dimensional surface traps are less efficient at generating quadrupolar potentials than three-dimensional geometries, higher RF voltages are needed to achieve comparable confinement strengths. For this reason, in all

subsequent simulations and experiments, both RF electrodes are driven at a frequency of $\Omega_{\mathrm{RF}} = 2\pi \times 19\,\mathrm{MHz}$, such that less voltage is needed for a target radial frequency, according to Equation 2.10 and Equation 2.12. Since the wavelength corresponding to this drive, $\lambda \approx 15\,\mathrm{m}$, is much larger than the physical length of the RF electrodes (a few millimetres), the spatial variation of the RF phase across the trap surface can be safely neglected.

Within the pseudopotential approximation, the total RF electric field at the ion position can be expressed as the superposition of the contributions from the inner and outer RF electrodes:

$$\vec{E}_{\mathrm{RF1}}(x,y,z,t) = \vec{E}_0^{(1)}(x,y,z)\,\cos(\Omega_{\mathrm{RF}}t)\,(\cos\alpha_{\mathrm{RF1}} + i\sin\alpha_{\mathrm{RF1}}), \tag{5.1a}$$

$$\vec{E}_{\mathrm{RF2}}(x,y,z,t) = \vec{E}_0^{(2)}(x,y,z)\,\cos(\Omega_{\mathrm{RF}}t)\,(\cos\alpha_{\mathrm{RF2}} + i\sin\alpha_{\mathrm{RF2}}), \tag{5.1b}$$

where $\vec{E}_0^{(1)}$, $\vec{E}_0^{(2)}$ are the field amplitudes produced by the inner and outer electrodes, respectively, and α_{RF1}, α_{RF2} are their corresponding phases. Since the proportionality between the electric field amplitude and the applied voltage holds also for rectangular electrode geometries [62], the field components can equivalently be written as:

$$\vec{E}_0^{(1)}(x,y,z) = V_{\mathrm{RF1}}\,\vec{E}_1(x,y,z), \tag{5.2a}$$

$$\vec{E}_0^{(2)}(x,y,z) = V_{\mathrm{RF2}}\,\vec{E}_2(x,y,z), \tag{5.2b}$$

where V_{RF1}, V_{RF2} are the adimensional voltage amplitudes applied to the inner and outer RF electrodes, while $\vec{E}_1(x,y,z)$, $\vec{E}_2(x,y,z)$ denote the normalized electric field distributions generated when a potential of 1 V is applied to each electrode individually.

The pseudopotential is then defined as:

$$\Phi_{\mathrm{PS}}(x,y,z) = \frac{Q^2}{2m\Omega_{\mathrm{RF}}^2}\left\langle \left|\vec{E}_{\mathrm{RF1}}(x,y,z,t) + \vec{E}_{\mathrm{RF2}}(x,y,z,t)\right|^2 \right\rangle_T, \tag{5.3}$$

where the brackets $\langle\cdot\rangle_T$ denote the temporal average over one RF cycle of period $T = 2\pi/\Omega_{\mathrm{RF}}$, as explained in Subsection 2.1.4.

From Equation 5.1a and Equation 5.1b, two control variables that govern the pseudopotential landscape for a fixed trap geometry can be identified: (i) the voltages applied to the RF electrodes, V_{RF1} and V_{RF2}, and (ii) the phases of the applied RF signals, α_{RF1} and α_{RF2}. It is therefore convenient to define two quantities that will be used throughout

this chapter:

$$\xi_{\mathrm{RF}} = \frac{V_{\mathrm{RF1}}}{V_{\mathrm{RF2}}}, \qquad \textit{RF amplitude ratio} \tag{5.4}$$

$$\Delta\alpha_{\mathrm{RF}} = \alpha_{\mathrm{RF1}} - \alpha_{\mathrm{RF2}}, \qquad \textit{RF phase mismatch} \tag{5.5}$$

In the following, we will examine how these parameters influence the pseudopotential landscape. For the simulations described in this section, the DC electrodes are set to ground, as only the RF pseudopotential is investigated.

Pseudopotential control: role of the RF amplitude ratio

The first aspect we investigate is how the RF amplitude ratio influences the pseudopotential landscape of the 4S trap. In all amplitude noise simulations, the outer RF electrodes are held at a fixed amplitude of 100 V, while the inner RF electrode amplitude is varied according to $V_{\mathrm{RF1}} = \xi_{\mathrm{RF}} V_{\mathrm{RF2}}$, with the chosen RF drive frequency ($\Omega_{\mathrm{RF}} = 2\pi \times 19$ MHz). Specifically, we consider RF amplitude ratios ξ_{RF} ranging from -1 to 1.5. For positive values of ξ_{RF}, both RF signals are kept in phase on the trap, i.e., $\alpha_{\mathrm{RF1}} = \alpha_{\mathrm{RF2}} = 0$. For negative values of ξ_{RF}, the RF signals are 180° out of phase. The pseudopotential is computed according to Equation 5.3 under these conditions.

Due to the underlying trap geometry, the RF pseudopotential exhibits two principal transport regimes, each characterized by the presence of two distinct minima and their corresponding saddle points. In the first configuration, referred to as the *vertical shuttling regime*, the two minima are aligned along the vertical axis. In the second con-

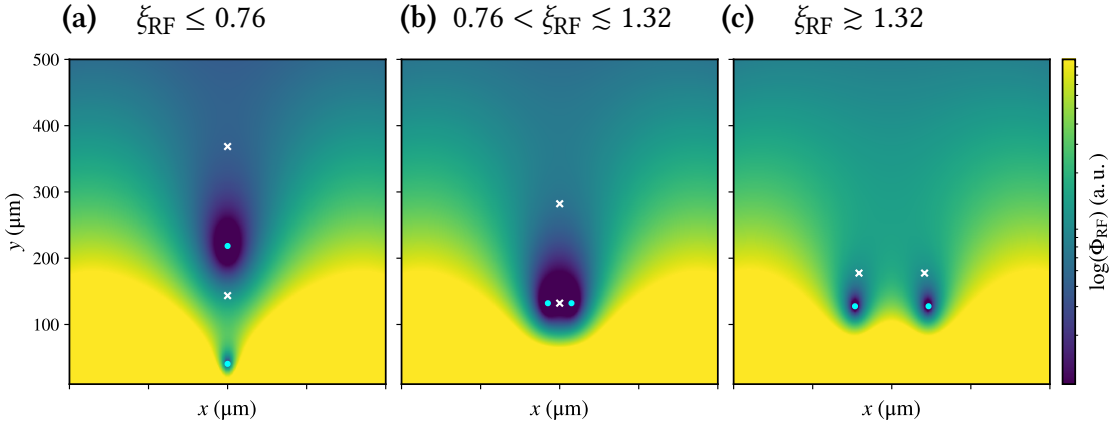

Figure 5.3: Pseudopotential landscape in the radial plane (xy). Cyan dots indicate the minima and white crosses the saddle points. The alignment of the minima shows the shuttling direction, vertical in panel (a) and horizontal in panels (b) and (c), while the orientation of the saddle points distinguishes the coupled (b) and uncoupled (c) horizontal transport regimes.

figuration, known as the *horizontal shuttling regime*, the minima instead align along the horizontal axis. This regime, relevant for the QSA architecture, can be further divided into two sub-regimes depending on the relative depth and separation of the wells, the *coupled* and *uncoupled* configurations. A radial cross-section of the pseudopotential in the double-well configuration is shown in Figure 5.3: panel (a) illustrates the vertical shuttling regime, panel (b) the coupled regime, and panel (c) the uncoupled regime.

Panel (a) and panel (b) of Figure 5.4 show the evolution of the positions of the pseudopotential minima as a function of the RF amplitude ratio ξ_{RF}. The three colour zone identify the three different regimes previously cited:

1. **Vertical shuttling ($\xi_{\mathrm{RF}} \leq 0.76$):** The two RF nulls move along the y-axis, approaching each other until they merge at $\xi_{\mathrm{RF}} \approx 0.76$. Simultaneous operation of the two linear traps is difficult in this regime, as the efficiency and q-factor of the minima differ significantly due to their unequal distances from the trap surface (see Subsection 2.1.2). While not directly relevant for the QSA architecture, this regime is important experimentally: ions are first trapped at $\xi_{\mathrm{RF}} = 0$ and then transported vertically to characterize both ξ_{RF} and $\Delta\alpha_{\mathrm{RF}}$. Beyond $\xi_{\mathrm{RF}} = 0.76$, the minima separate horizontally, marking the transition to the horizontal shuttling regime.

2. **Horizontal shuttling ($\xi_{\mathrm{RF}} > 0.76$):** The two minima move along the x-axis. In the coupled regime the RF nulls are sufficiently close that their Coulomb coupling is non-negligible, while in the uncoupled regime they are well separated. The saddle points are vertically aligned in the coupled regime, while they are aligned horizontally in the uncoupled one (see Figure 5.3):

   (a) *Coupled regime ($\xi_{\mathrm{RF}} \lesssim 1.32$):* the pseudopotential forms a double-well, with one saddle point defining the barrier between wells and the other one the trap depth.

   (b) *Uncoupled regime ($\xi_{\mathrm{RF}} \gtrsim 1.32$):* as ξ_{RF} increases, both the well separation and barrier height grow until the saddle points align horizontally. In this regime, the Coulomb coupling between separate wells is negligible, and information can be stored.

The radial secular frequencies as a function of the RF amplitude ratio ξ_{RF} are shown in panel (c) of Figure 5.3. During vertical transport, the radial frequencies at the two pseudopotential minima differ substantially, which renders the simultaneous operation of both linear traps impractical, except near the transition from vertical to horizontal shuttling. At this transition, the radial frequencies vanish, as the quadratic component of the pseudopotential goes to zero. In the horizontal regime, the radial frequencies become identical due to the trap symmetry and increase monotonically with ξ_{RF}.

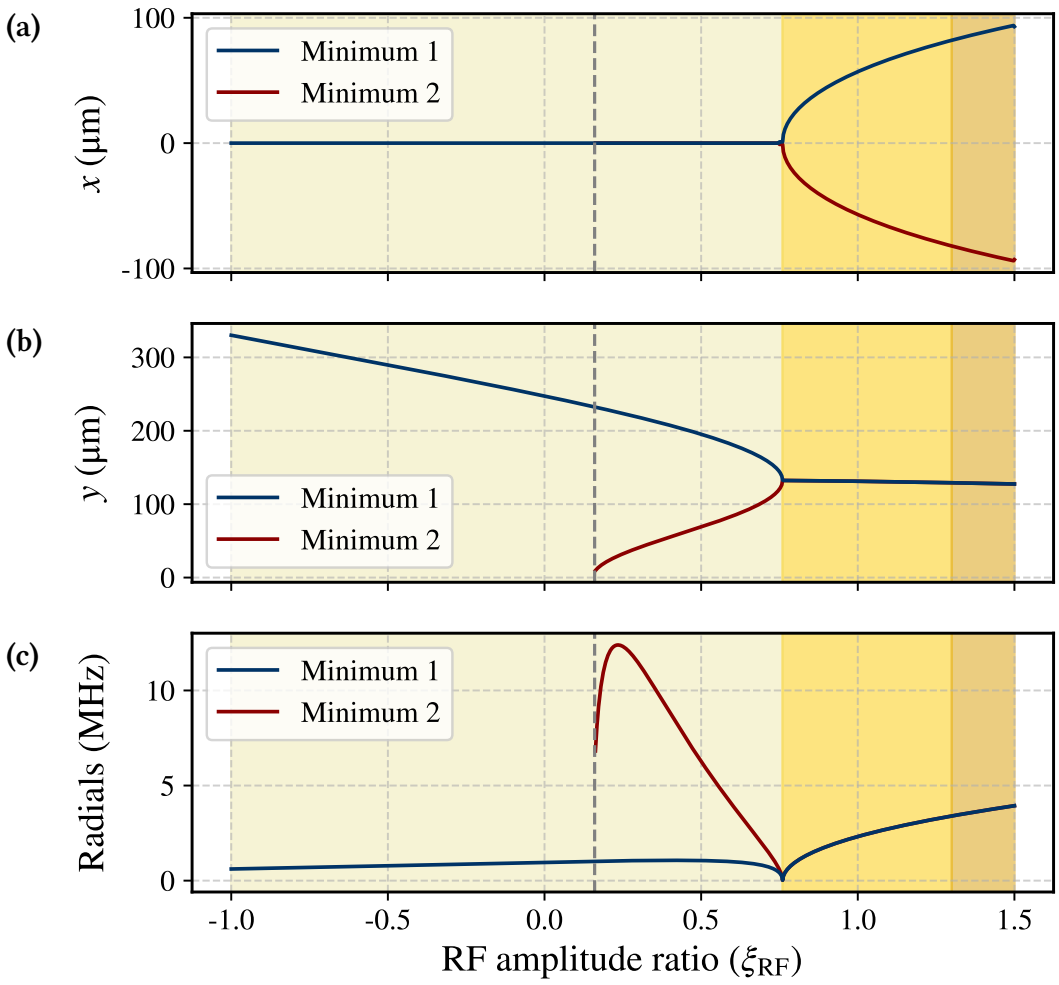

Figure 5.4: Evolution of the RF nulls positions and radial secular frequencies vs. ξ_{RF}. Coloured zones indicate different regimes (from left to right, respectively, vertical, coupled, uncoupled), and the dashed line marks the appearance of the second minimum above the trap surface. (a) Horizontal x positions, (b) vertical y positions of the RF nulls, (c) radial trapping frequencies.

Pseudopotential control: role of the RF phase mismatch

We now discuss the effects of a phase mismatch between the two RF signals applied to the trap on the pseudopotential landscape. The pseudopotential, according to Equation 5.3, depends on the squared magnitude of the sum of the electric fields generated by the RF1 and RF2 electrodes, which in turn are determined by the phases of the two RF signals. In all phase mismatch simulations, the amplitudes of both RF1 and RF2 electrodes are held fixed at 100 V, i.e. $\xi_{RF} = 1$, while the phase mismatch $\Delta\alpha_{RF}$ is varied, keeping the RF drive frequency constant at $\Omega_{RF} = 2\pi \times 19$ MHz. Specifically, we consider phase mismatches $\Delta\alpha_{RF}$ ranging from $-\pi/2$ to $\pi/2$.

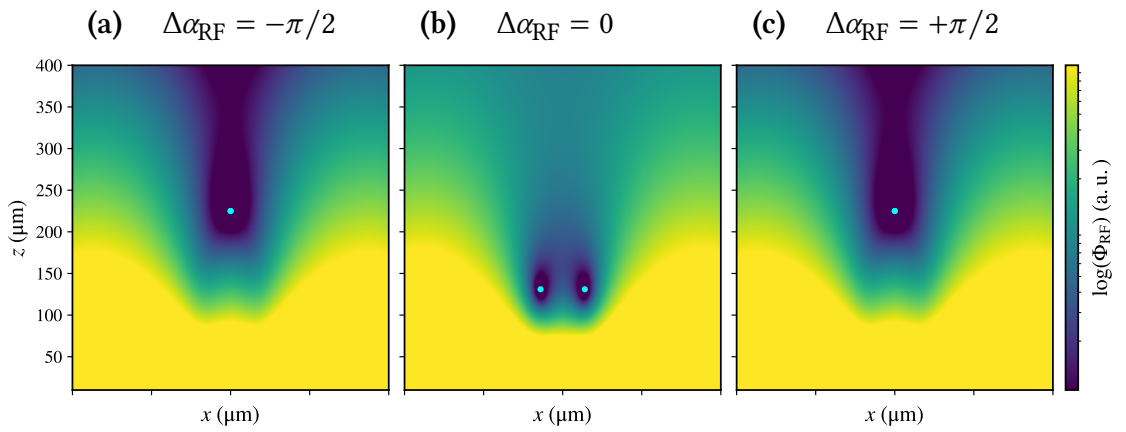

FIGURE 5.5: Variation of the positions of the pseudopotential minima with RF phase mismatch $\Delta\alpha_{RF}$: the double-well structure is fully developed at $\Delta\alpha_{RF} = 0$ (b) and gradually disappears for nonzero phase differences (a) and (c).

The first notable effect of varying the RF phase mismatch is the displacement of the pseudopotential minima, showing a trend similar to that observed when varying the RF amplitude ratio (see Figure 5.3). As discussed above, the condition $\xi_{RF} = 1$ corresponds to the coupled regime, characterized by the typical double-well potential formed by the saddle points. Figure 5.5 illustrates three representative phase differences: panel (a) $\Delta\alpha_{RF} = -\pi/2$, panel (b) $\Delta\alpha_{RF} = 0$, and panel (c) $\Delta\alpha_{RF} = +\pi/2$. It is evident that the pseudopotential landscape is drastically modified if the applied RF signals are not in phase. Consequently, precise control of the relative RF phase is essential to preserve the coupled regime and the associated double-well configuration.

The second important consequence of a phase mismatch is that the pseudopotential minimum is no longer a RF null. This implies the presence of a residual electric field at the ion's position, which drives excess micromotion and increases the ion's kinetic energy, leading to motional heating.

5.3 RF TRANSPORT: HEATING RATE SIMULATIONS

From the analysis presented in the previous sections, we have seen that variations in both the RF amplitude ratio and the relative phase of the RF signals alter the pseudopotential landscape. Thus, any source of technical noise in the RF circuitry, affecting either the amplitude or the phase of the electric fields generated by the electrodes, translates directly into fluctuations of the pseudopotential, leading to excess micromotion and ion heating. In this section, using the model discussed in Chapter 3, we simulate ion heating during RF transport in the 4S trap.

5.3.1 Amplitude noise model

In Subsection 3.2.1, we developed a model to quantify the heating rates induced by technical amplitude noise. In the specific case of the 4S trap, unlike the junction configuration, two independent RF signals are applied; this requires adopting the extended model that accounts for multiple RF drives.

Considering Equation 5.1a and Equation 5.1b, which describe the individual electric field contributions from the inner and outer RF electrodes, we fix the outer RF voltage amplitude to 100 V, while varying the inner one according to $V_{\mathrm{RF1}} = \xi_{\mathrm{RF}}\, V_{\mathrm{RF2}}$. Amplitude noise is then injected on both RF components, yielding the total electric field, obtained from Equation 3.23, as:

$$\begin{aligned}\vec{E}_{\mathrm{RF}}(x,y,z,t) = \vec{E}_0^{(1)}(x,y,z)\left[\cos(\Omega_{\mathrm{RF}}t) + \chi_{\mathrm{N}}^{(1)}\cos\Big((\Omega_{\mathrm{RF}}+\omega)t + \varphi^{(1)}(t)\Big)\right] + \\ + \vec{E}_0^{(2)}(x,y,z)\left[\cos(\Omega_{\mathrm{RF}}t) + \chi_{\mathrm{N}}^{(2)}\cos\Big((\Omega_{\mathrm{RF}}+\omega)t + \varphi^{(2)}(t)\Big)\right],\end{aligned} \tag{5.6}$$

where $\varphi^{(1)}(t)$ and $\varphi^{(2)}(t)$ are stochastic phase fluctuations that account for possible correlations between the noise contributions of the two RF sources. Following the derivation presented in Subsection 3.2.1, we treat these fluctuations as random processes, such that the noise contributions are completely uncorrelated. This relative dephasing between the two RF sources prevents complete noise cancellation, allowing for heating even when the ions are confined at the RF nulls, as described in Subsection 3.2.1.

In this case, the noise force contribution expressed in Equation 3.28 becomes:

$$\begin{aligned}\vec{F}_N = \frac{Q^2}{4m\Omega_{\mathrm{RF}}^2}\left[\vec{\nabla}\Big(\vec{E}_0^{(1)}\cdot\vec{E}_0^{(2)}\Big) + \vec{\nabla}\Big(\vec{E}_0^{(1)}\Big)^2\right]\chi_N^{(1)}\cos\Big(\omega t + \varphi^{(1)}(t)\Big) + \\ + \frac{Q^2}{4m\Omega_{\mathrm{RF}}^2}\left[\vec{\nabla}\Big(\vec{E}_0^{(2)}\cdot\vec{E}_0^{(1)}\Big) + \vec{\nabla}\Big(\vec{E}_0^{(2)}\Big)^2\right]\chi_N^{(2)}\cos\Big(\omega t + \varphi^{(2)}(t)\Big).\end{aligned} \tag{5.7}$$

Finally, the total heating rate can be directly derived from Equation 3.30:

$$\begin{aligned}\dot{\bar{n}} = \frac{Q^4}{16m^3\Omega_{\mathrm{RF}}^4\hbar\omega}\left[\vec{\nabla}\Big(\vec{E}_0^{(1)}\cdot\vec{E}_0^{(2)}\Big) + \vec{\nabla}\Big(\vec{E}_0^{(1)}\Big)^2\right]^2 \frac{S_{V_N}^{(1)}(\Omega_{\mathrm{RF}}\pm\omega)}{\Big(V_0^{(1)}\Big)^2} + \\ + \frac{Q^4}{16m^3\Omega_{\mathrm{RF}}^4\hbar\omega}\left[\vec{\nabla}\Big(\vec{E}_0^{(2)}\cdot\vec{E}_0^{(1)}\Big) + \vec{\nabla}\Big(\vec{E}_0^{(2)}\Big)^2\right]^2 \frac{S_{V_N}^{(2)}(\Omega_{\mathrm{RF}}\pm\omega)}{\Big(V_0^{(2)}\Big)^2}).\end{aligned} \tag{5.8}$$

5.3.2 Phase noise model

We now turn to simulations that include technical phase noise, using the model developed in Subsection 3.2.2. As noted previously, in the case of the 4S trap, unlike the junction configuration, two independent RF signals are applied. Consequently, a phase-noise-induced heating rate is indeed present, as we must adopt the extended model that accounts for multiple RF drives. In this scenario, phase noise is injected on both RF components, so that the total electric field, derived from Equation 3.23, reads as:

$$\begin{aligned}\vec{E}_{\mathrm{RF}}(x,y,z,t) =& \vec{E}_0^{(1)}(x,y,z)\cos\left(\Omega_{\mathrm{RF}}t+\theta_N^{(1)}(t)\right)+\\ &+\vec{E}_0^{(2)}(x,y,z)\cos\left(\Omega_{\mathrm{RF}}t+\theta_N^{(2)}(t)\right).\end{aligned} \tag{5.9}$$

Following the derivation presented in Subsection 3.2.2, the noise force contribution given in Equation 3.41 becomes:

$$\vec{F}_N = -\frac{Q^2}{4m\Omega_{\mathrm{RF}}^2}\left(\vec{\nabla}\left[\vec{E}_0^{(1)}\cdot\vec{E}_0^{(2)}\right]+\vec{\nabla}\left[\vec{E}_0^{(2)}\cdot\vec{E}_0^{(1)}\right]\right)\chi_N^2\cos(2\omega_N t)\,, \tag{5.10}$$

where $\chi_{12}=\chi_{21}=\chi_N$ represents the amplitude of the combined noise on the two RF signals, as defined in Equation 3.39. This is appropriate because the physically relevant quantity is the *phase difference* between the two signals, allowing us to characterized the effect using a single single effective noise amplitude and a single spectral phase noise density $S_{\Theta_N}\left(\Omega_{\mathrm{RF}}+\frac{\omega}{2}\right)$. Finally, the heating rate follows directly from Equation 3.43:

$$\dot{\bar{n}} = \frac{Q^4}{16m^3\Omega_{\mathrm{RF}}^4\hbar\omega}\left[\left(\vec{\nabla}\left[\vec{E}_0^{(1)}\cdot\vec{E}_0^{(2)}\right]+\vec{\nabla}\left[\vec{E}_0^{(2)}\cdot\vec{E}_0^{(1)}\right]\right)^2 S_{\Theta_N}\left(\Omega_{\mathrm{RF}}+\frac{\omega}{2}\right)\right]. \tag{5.11}$$

5.3.3 Heating rates simulations

To quantitatively assess the impact of technical RF noise on the ion's motional excitation, we compute the heating rates using the formalism derived in the previous sections. All geometric and electrical parameters entering Equation 5.8 and Equation 5.11 are obtained from the trap designs discussed in Subsection 5.2.1. The heating rates are evaluated as a function of the RF amplitude ratio ξ_{RF}, using the coordinate system defined in Figure 5.2. In all plots, the radial secular frequency is shown for comparison, and the light yellow region highlights the interval of ξ_{RF} corresponding to horizontal shuttling.

Amplitude noise

An estimate of the ion heating due to uncorrelated amplitude noise can be obtained by introducing two assumptions about the spectral voltage noise density evaluated at the relevant motional frequency. First, we assume that the noise spectral density of each source depends on its output voltage. Second, we assume that two RF sources operated at the same output voltage exhibit the same noise spectral density. Under these assumptions, one can write:

$$\frac{S_{V_N}^{(1)}(\Omega_{\mathrm{RF}} \pm \omega)}{\left(V_0^{(1)}\right)^2} \approx \frac{S_{V_N}^{(2)}(\Omega_{\mathrm{RF}} \pm \omega)}{\left(V_0^{(2)}\right)^2} = \bar{S}_{V_N}(\Omega_{\mathrm{RF}} \pm \omega), \tag{5.12}$$

These assumptions are justified if we consider that the two sources are equal, and if each RF generator is modelled as a fixed voltage source whose output is regulated with an attenuator, such that the amount of noise at the output of the source is also defined by the state of the attenuator.

In the absence of an explicit circuit model, we assume a reference value for the spectral voltage noise density typical of white noise in well-stabilized laboratory RF sources [44, 66], for a fixed output voltage of 100 V:

$$S_{V_N}(\Omega_{\mathrm{RF}} \pm \omega) \sim 1 \times 10^{-14}\,\mathrm{V^2/Hz}\,. \tag{5.13}$$

This value scales with the output voltage as ξ_{RF}^2, as expressed in Equation 3.22.

Using Equation 5.8, we compute the heating rates for all three motional modes in the case of amplitude noise. The results are reported in Figure 5.6. As expected, the axial heating rate is zero for all ξ_{RF}, since the RF field and its gradient vanish along the trap axis, because no axial confinement is provided. The discussion therefore focuses on the radial modes.

The heating rate along the x-axis, $\dot{\bar{n}}_x$, is negligible throughout the vertical shuttling regime, as the RF minima move solely along y (see Figure 5.4) and the field gradient along x remains zero by symmetry. When the system enters the horizontal shuttling regime, the two RF nulls begin to separate along x, and $\dot{\bar{n}}_x$ increases monotonically, reaching a maximum of $\dot{\bar{n}}_x \approx 1.53\times10^5$ phonon/s. This value is significantly larger than the 10^1–10^3 phonon/s typically observed in linear Paul traps with similar RF-source noise levels [96, 97], indicating strong motional excitation in this regime.

On the other hand, the heating rate along the y-axis, $\dot{\bar{n}}_y$, exhibits a different behaviour.

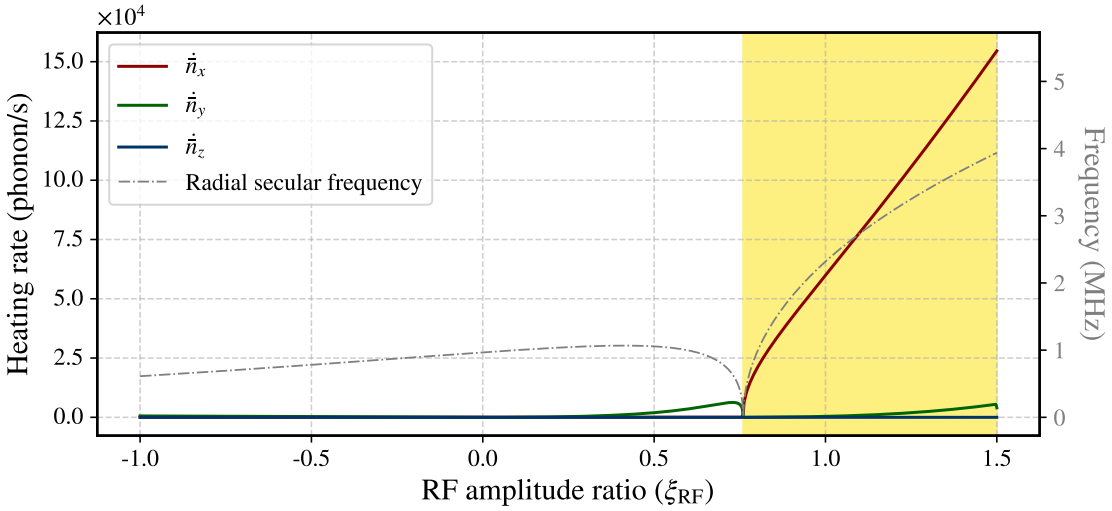

FIGURE 5.6: Heating rates induced by amplitude noise as a function of the RF amplitude ratio ξ_{RF}. The radial secular frequencies are shown for reference. The shaded yellow region marks the horizontal shuttling regime.

During vertical shuttling, heating along y is always present and slowly increases as the two minima approach each other, reaching at most $\dot{\bar{n}}_y \approx 7.5 \times 10^3$ phonon/s. Upon entering the vertical regime, it drops to zero, and then increases again in the horizontal regime due to a small vertical displacement of the RF nulls (on the order of a few micrometres), resulting in a residual nonzero gradient along y. Overall, $\dot{\bar{n}}_y$ remains roughly thirty times smaller than $\dot{\bar{n}}_x$ during the horizontal shuttling regime.

PHASE NOISE

In the case of phase noise, a single spectral phase noise density is present (see Equation 5.11). As before, we consider a reference value typical of white noise in well-stabilized laboratory RF sources [44, 66]:

$$S_{\Theta_N}\left(\Omega_{RF} \pm \frac{\omega}{2}\right) \sim 1 \times 10^{-14}/\text{Hz} \,. \tag{5.14}$$

This time this value doesn't depend on the output voltage of the attenuator, since attenuation affects only the signal amplitude and not its phase.

Using Equation 5.11, we compute the heating rates for all three motional modes in the case of phase noise. The results are shown in Figure 5.7. As in the amplitude-noise case, the axial mode does not heat up, and the relevant dynamics occur only in the radial directions.

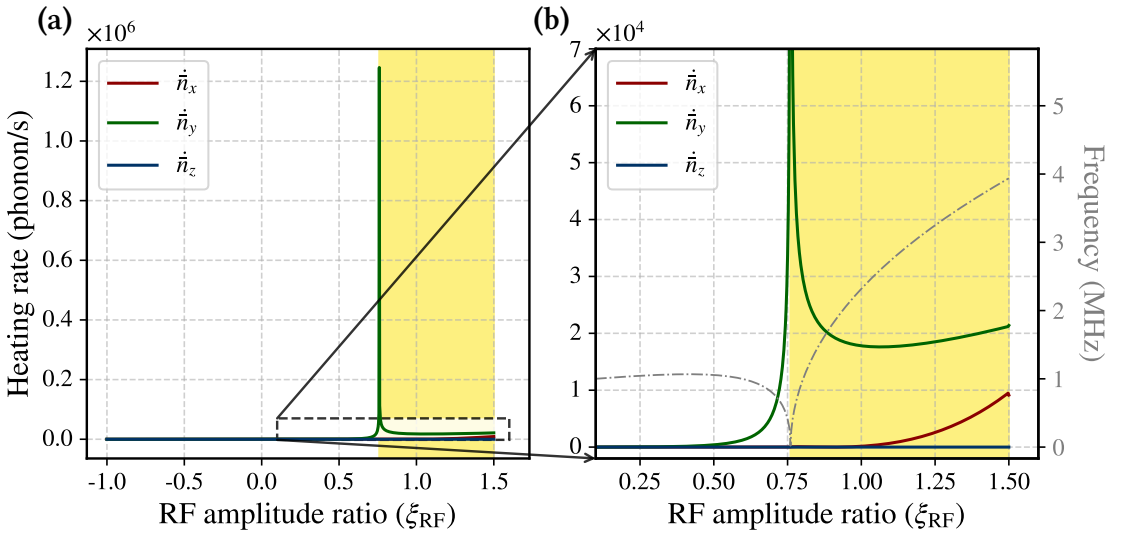

FIGURE 5.7: Heating rates induced by phase noise as a function of the RF amplitude ratio ξ_{RF}. (a) Heating rates for the axial and radial modes. (b) Zoom on the area of interest, with the radial secular frequency shown for reference. The shaded yellow region marks the horizontal shuttling regime.

The heating rate along the x-axis, $\dot{\bar{n}}_x$, mirrors the structural behaviour seen for amplitude noise: it is negligible in the vertical regime, then increases monotonically during horizontal shuttling as the RF nulls separate along x, even though with a different behaviour than in the amplitude-noise case. It reaches a peak value of $\dot{\bar{n}}_x \approx 0.9 \times 10^4$ phonon/s, which is about an order of magnitude smaller than in the amplitude-noise case, yet still significantly larger than typical heating rates in standard Paul traps [96, 97]. Thus, even in the phase-noise scenario, the ion experiences substantial motional excitation along the x-axis during horizontal shuttling.

The heating rate along the y-axis, $\dot{\bar{n}}_y$, exhibits a behaviour which differs from the amplitude noise case. It is small but nonzero in the vertical regime, and diverges near the transition from vertical to horizontal shuttling due to the vanishing radial secular frequency (see panel (b) of Figure 5.7). After this divergence, $\dot{\bar{n}}_y$ decreases again and settles to values around $\dot{\bar{n}}_y \approx 2.1 \times 10^4$ phonon/s.

Discussion

In both the amplitude and phase noise scenarios, the simulations indicate that the heating rates during vertical shuttling are orders of magnitude smaller compared to the horizontal shuttling regime. More specifically, for amplitude noise, the heating is primarily along the x-axis, with $\dot{\bar{n}}_y$ remaining substantially smaller than $\dot{\bar{n}}_x$. For phase noise, the

situation is reversed, although the difference between the two radial modes is less pronounced. These contrasting behaviours reflect the different dependency of ion heating on the RF field gradients in the two cases.

It is important to note that the results presented here are based on the simplified assumption that both RF electrodes are subject to equal white noise. Therefore, the conclusions regarding the relative heating rates along the radial directions are strictly valid only within this context. More general scenarios could involve unequal noise amplitudes on the two electrodes, voltage-dependent noise, or noise applied selectively to either the inner or the outer electrodes. Investigating these alternative configurations would provide a more complete characterization of the influence of technical noise on the ion's motional excitation.

5.4 Comparison with experimental data

In this final section of the chapter, we compare the predictions of our previously developed models with experimental measurements. We first provide a brief overview of the experimental setup, then describe the heating rate measurements, and, finally, assess how well our models reproduce the observed data.

5.4.1 Experimental setup

The experiments presented in the following section were performed using a cryogenic ion-trapping setup designed for $^{40}Ca^{+}$ ions. Detailed information on the complete experimental system used for the RF transport experiment can be found in the PhD theses of Michael Niedermayr [98], Philip Holz [95], and Marco Valentini (under preparation).

A close-up of the experimental setup is shown in panel (a) of Figure 5.8. The cryostat is a closed-cycle Gifford-MacMahon cooler operating at a base temperature of approximately 8 K. The trap itself is mounted within an octagon vacuum chamber in a socket, consisting of three main elements: a U-shaped mechanical support, a filter PCB, and a trap carrier PCB, which together provide the necessary electrical connections, mechanical stability, and cooling for the trap. Two RF resonators, located within the 8 K shield and driven by two independent phase-locked RF sources, amplify and filter the signals required for radial confinement of ions. Finally, the optical access for the laser beams is provided by viewports in the octagon vacuum chamber. Ions are imaged through an additional window at the bottom flange of the setup. A cross-section of the octagon showing the laser beam paths is presented in panel (b) of Figure 5.8. The Doppler cool-

ing, repumping, and photoionization lasers are superimposed and enter the chamber from the same input viewport, while the qubit laser (729 nm) follows a second, perpendicular path. Ions are loaded by photoionization of a neutral calcium flux from an oven mounted in a dedicated feed-through. Two pairs of Helmholtz coils attached outside the chamber generate a magnetic field at 45° relative to all laser beams, defining the ions' quantization axis.

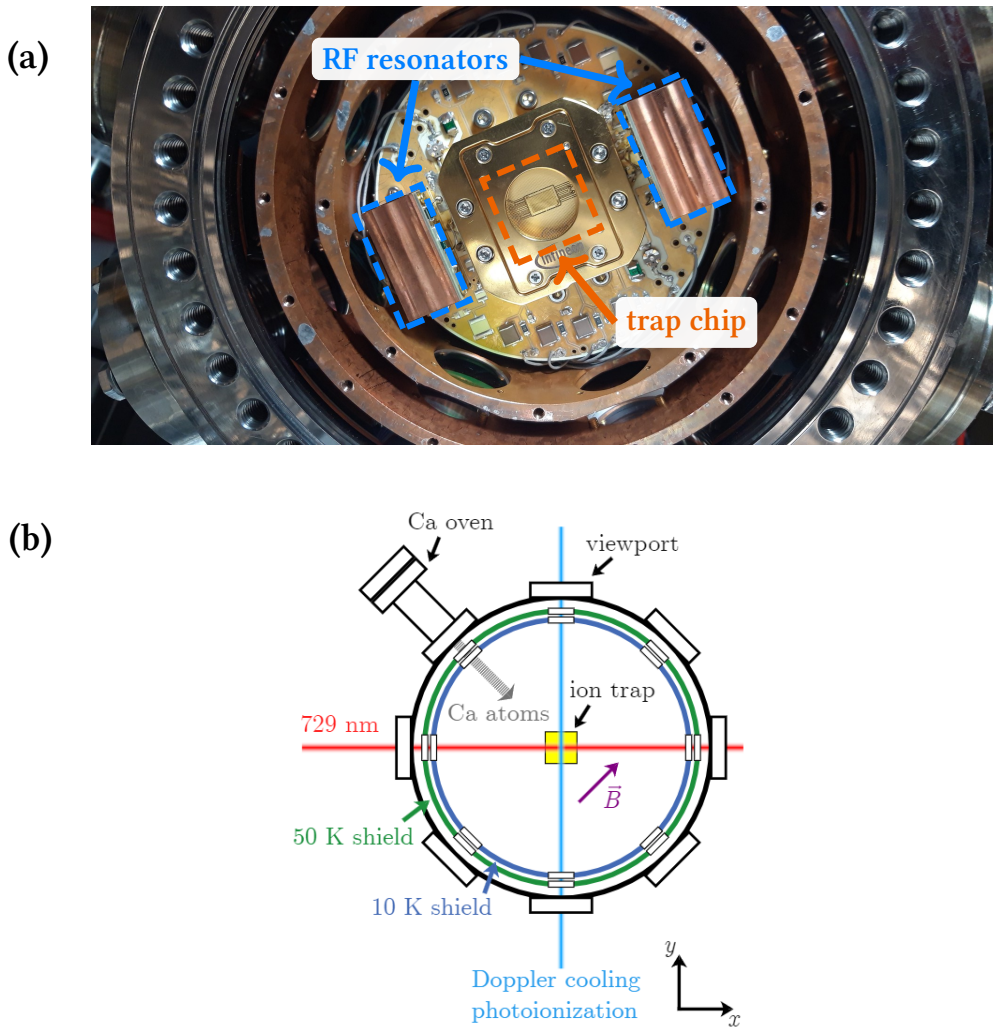

FIGURE 5.8: Experimental setup. (a) The cryostat provides a two-stage cooling environment (50 K intermediate shield and 8 K cold finger) for the ion trap. The trap chip (orange) is mounted in a socket with a U-shaped support, filter PCB, and carrier PCB, while the RF resonator circuits (blue) deliver the radial confinement signals at 10–20 MHz. (b) Optical access of the experimental setup. Viewports provide the laser path for Doppler cooling, repumping, and photoionization, with a perpendicular path for the 729 nm qubit laser. The trap axis in this schematic should be oriented towards the atomic oven (indicated by a 45° tilt of the yellow square). Image courtesy of Marco Valentini.

5.4.2 Heating rate measurements

The heating rate data presented in this work were measured by Marco Valentini and Martin van Mourik on the 4S trap using both RF resonators. The heating rates were estimated by measuring the decay of carrier Rabi flops at different waiting times, which provides an effective measure of the overall motional excitation, although without distinguishing between axial and radial modes [44]. Panel (a) in Figure 5.9 shows the heating-rate measurements on a linear scale, together with the radial secular frequencies encoded as a colorbar. Panel (b) presents the same dataset but with a logarithmic vertical axis. This log-scale view is useful both to highlight trends that are not easily discernible in linear scale and to accommodate the large spread in observed heating rates, which differ by more than an order of magnitude. In the vertical shuttling regime, heating rates remain relatively low compared to the horizontal regime, but they increase sharply as ξ_{RF} is raised. In particular, for $\xi_{\mathrm{RF}} > 1$, corresponding to horizontal radial transport above the trap surface, heating rates exceed $\dot{n} > 40$ phonons/ms. This behavior is not consistent with simple surface-noise models, which predict a strong dependence on the ion-surface distance d (typically $\dot{n} \propto d^{-4}$ [40]); in the horizontal transport regime, d remains nearly constant. Furthermore, the heating rates observed with the dual-resonator configuration are several orders of magnitude higher than those reported for a single-resonator setup at $\xi_{\mathrm{RF}} = 1$, measured to be less than 10 phonon/s for axial frequencies ranging from around 0.3 MHz to 1.2 MHz.

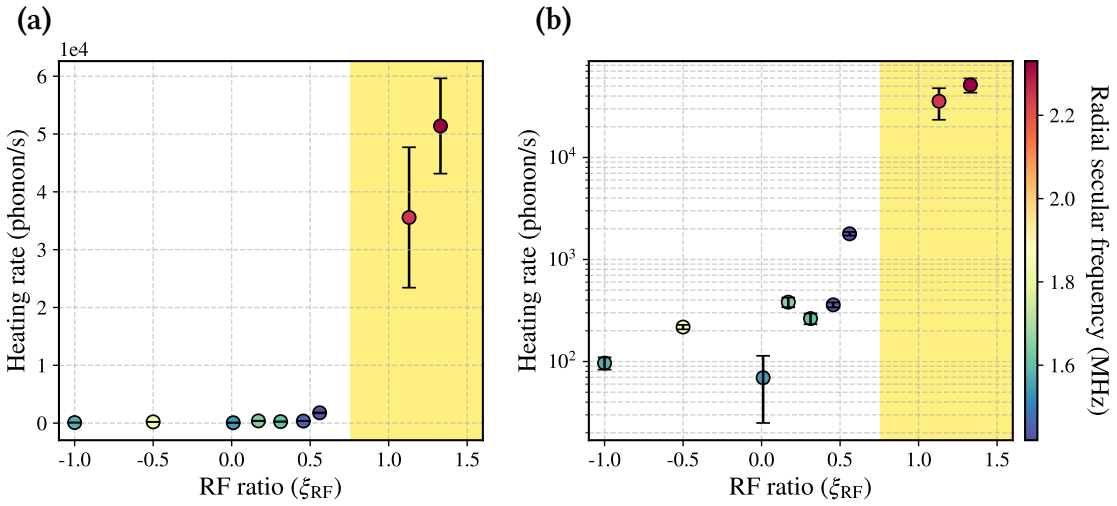

Figure 5.9: Heating rates measured on the 4S trap using carrier Rabi flops at different waiting times. The heating rates increases sharply for $\xi_{\mathrm{RF}} > 1$, where radial transport occurs along the trap surface (yellow region). Panel (a) shows the data on a linear scale, while panel (b) uses a logarithmic scale to better visualize the large spread in heating rates.

A more detailed analysis of the heating mechanisms would require measurements on the axial and radial sidebands, which would allow one to investigate the specific ion heating for each mode. However, the measurements were not performed as it was not possible to cool the ions down to the resolved sideband regime due to the high heating rate observed.

5.4.3 Simulations–measurements comparison

We now compare the models derived in Section 5.3 with the measured heating rates presented in the previous section. Before presenting the results, we note that the experimentally measured secular frequencies (see Figure 5.9) differ from those shown in Figure 5.4. This discrepancy arises primarily because the effective trap parameters used in the experiment, such as the RF amplitude, differ slightly from those assumed in the model. To account for this, we rescale the experimental data so that the radial secular frequencies match those used in our simulations.

We focus only on the radial modes, as the axial heating rate, $\dot{\bar{n}}_z$, is negligible due to the absence of axial confinement. All relevant trap parameters in the model, including the RF drive amplitude Ω_{RF}, the radial secular frequency ω, the ion charge Q, and the mass m, are fixed. The only free parameter in the fit is the spectral noise density. Since the experiment does not provide direct information about the direction from which the noise couples to the ion, we cannot assume *a priori* whether the heating originates predominantly along the x or y radial directions. Therefore, both directions must be analysed independently.

To determine its value, we fit the theoretical heating-rate model to the experimental measurements using a non-linear least-squares procedure. In particular, we employ the CURVE_FIT routine from the SCIPY.OPTIMIZE package, which minimizes the weighted squared residuals between model and data while accounting for the experimental uncertainties. This yields an estimate of the spectral noise density together with the associated standard error, which is then compared to the reference value assumed for the spectral noise density (see Equation 5.13 and Equation 5.14).

The fitted values of the spectral noise density, along the two radial directions, for both amplitude and phase noise are reported in Table 5.1. These numerical results, however, should be considered together with the corresponding graphical comparisons between model and experiment, since, in some cases, the model reproduces only a subset of the measured data. In such situations, the fitted spectral noise density reflects primarily the region where the model provides a meaningful match, rather than the entire dataset.

	x-axis	y-axis
$S_{V_N}(\Omega_{\mathrm{RF}} \pm \omega)$	$(1.44 \pm 0.21) \times 10^{-14}\,\mathrm{V^2/Hz}$	$(2.44 \pm 0.07) \times 10^{-14}\,\mathrm{V^2/Hz}$
$S_{\Theta_N}\left(\Omega_{\mathrm{RF}} \pm \frac{\omega}{2}\right)$	$(8.49 \pm 1.29) \times 10^{-14}\,\mathrm{Hz^{-1}}$	$(2.70 \pm 0.08) \times 10^{-16}\,\mathrm{Hz^{-1}}$

Table 5.1: Spectral noise densities extracted for the X and Y axes from the fit of the model to the experimental heating rates. The parameters are obtained using a non-linear least squares method (CURVE_FIT) and are reported as mean ± standard error.

The amplitude noise comparison is shown in Figure 5.10. Panels (a) and (b) display the model and experimental data for the x-axis in linear and logarithmic scale, respectively, while panels (c) and (d) show the corresponding results for the y-axis. Both scales are included because they highlight different features of the agreement between model and experiment. From the linear plots, one might initially conclude that the model for $\dot{\bar{n}}_x$ provides a reasonable description across the full dataset. However, the logarithmic scale reveals that the fit along x actually follows only the points in the horizontal shuttling regime, while it fails to reproduce the behaviour at lower RF ratios. Consequently, the fitted spectral density for the x-axis carries a large uncertainty (the relative error is around 15%), which is expected given that the fit is effectively constrained by only two data points.

In contrast, the fit along the y-axis captures a larger fraction of the dataset and yields a spectral density that is compatible with the reference value. Although the model underestimates the sharp increase in $\dot{\bar{n}}_y$ observed during horizontal transport, the overall agreement is significantly better than in for $\dot{\bar{n}}_x$. This behaviour is qualitatively consistent with the expectation that the ion heats up more along the direction of motion during transport. We observe a critical value for $\xi_{\mathrm{RF}} = 0$, where the heating rate decreases to very low values. This behaviour is expected, as the system transitions into the single-drive configuration, where interference terms arising from the superposition of multiple RF drives are suppressed, resulting in a reduced noise contribution.

However, it is important to recall that the experimental data were acquired with the cooling and probing laser parallel to the trap surface (see panel (b) of Figure 5.8). In this configuration, the measurement is sensitive only to motional modes that have a non-zero projection along the laser wavevector $\vec{k}$, i.e. the axial and x-radial modes. For a laser propagating parallel to the trap surface, the component of $\vec{k}$ perpendicular to the surface is zero, and therefore there is no coupling to the y-radial mode. As a result, agreement with a model describing heating purely along the y-axis should not be expected, since motion parallel to the laser axis is effectively invisible to the measurement.

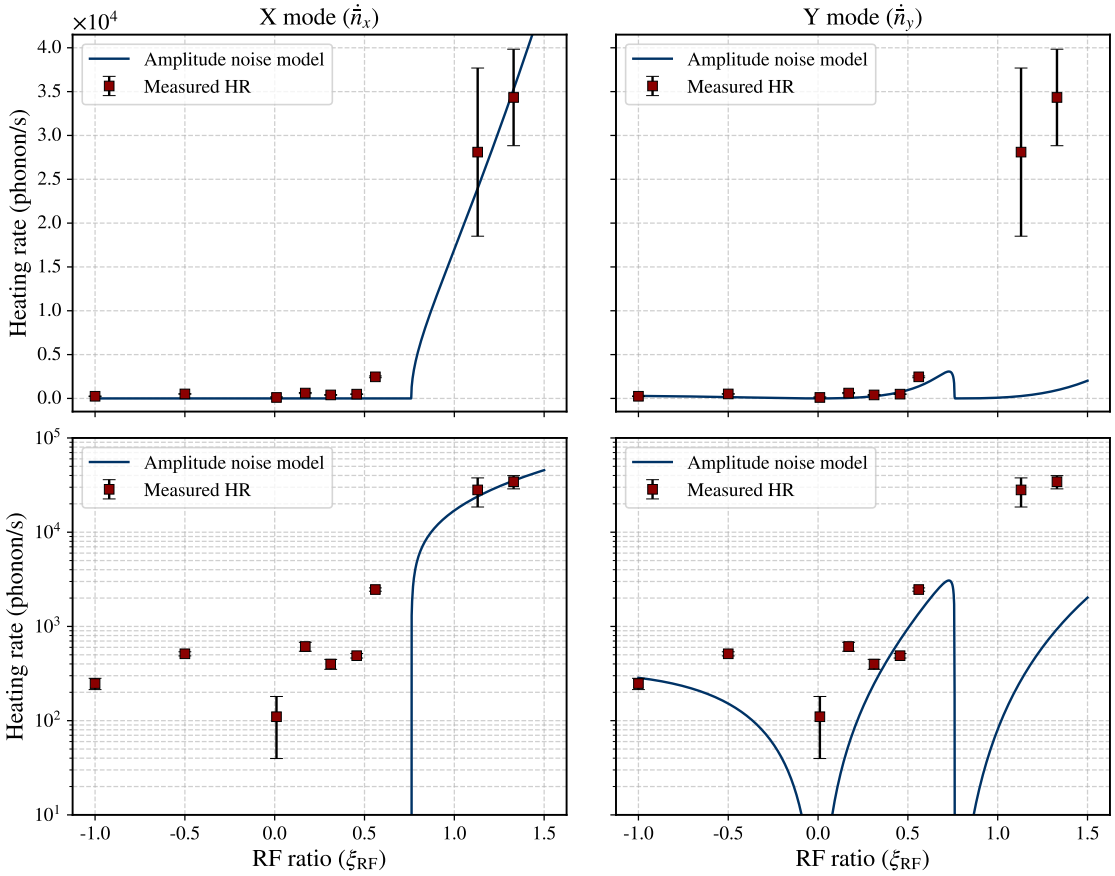

Figure 5.10: Comparison between model and experiment for amplitude–noise induced heating along the radial x and y directions. Panels (a,b) show the x–axis data and model in linear and logarithmic scale; panels (c,d) show the same for the y-axis. The fits illustrate that the y–axis model captures a larger portion of the dataset, while the x–axis fit is weakly constrained.

The partial agreement observed in the y-axis model thus indicates that the correspondence between the nominal trap axes and the actual motional directions may not be exact, suggesting the presence of additional asymmetries or mode rotations in the trap.

Moving to the phase-noise analysis, the comparison between model and experiment is shown in Figure 5.11. Panels (a) and (b) display the results for the x-axis in linear and logarithmic scale, respectively, while panels (c) and (d) show the corresponding plots for the y-axis. As in the amplitude-noise case, the x-axis fit reproduces only the data in the horizontal shuttling regime and fails to describe the whole dataset. The modelled heating rate even diverges for $\xi_{RF} = 1$, and the fitted spectral phase noise density is one order of magnitude larger than in the amplitude noise analysis, indicating a poor and weakly constrained fit.

The y–axis comparison instead shows substantially better qualitative agreement. Al-

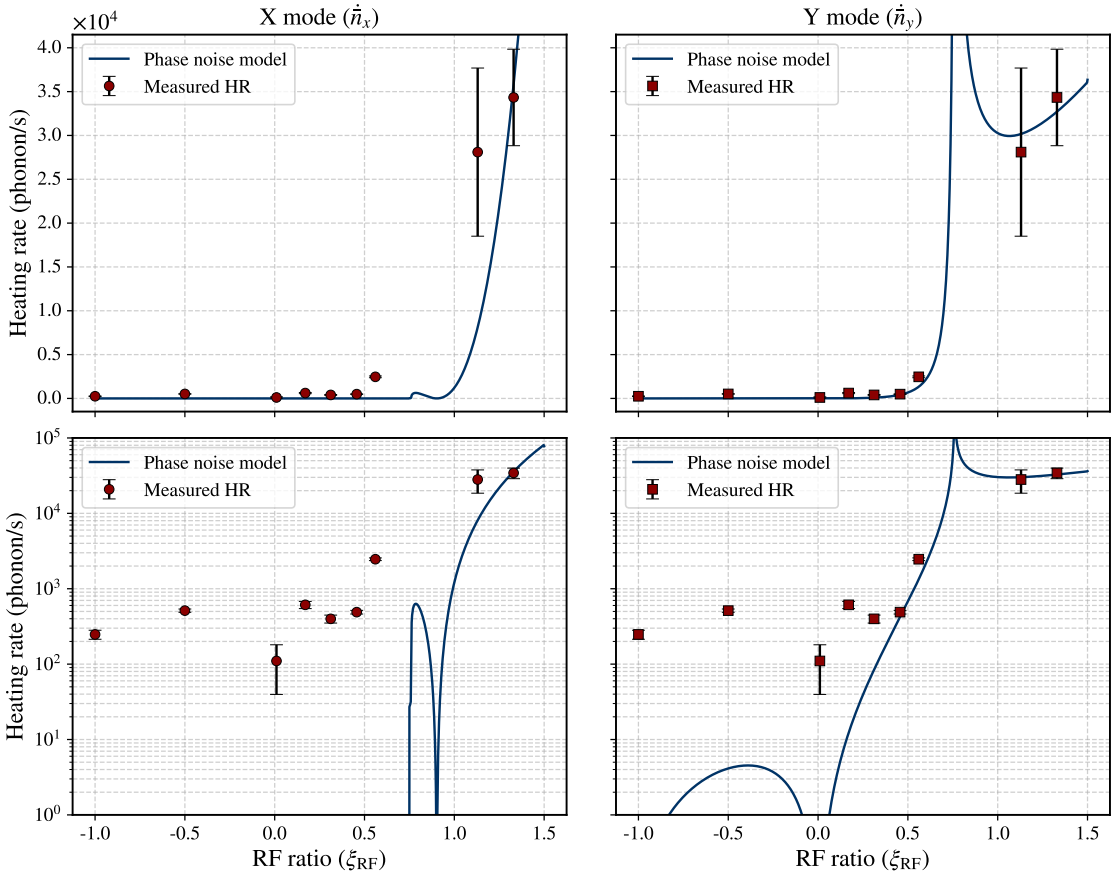

Figure 5.11: Comparison between model and experiment for phase–noise induced heating along the radial x and y directions. Panels (a,b) show the x–axis data and model in linear and logarithmic scale; panels (c,d) show the same for the y-axis. The fits illustrate that the y–axis model captures a larger portion of the dataset, while the x–axis fit is weakly constrained.

ready on the linear scale, the model appears capable of reproducing the overall trend of the data, unlike in the amplitude–noise case. The logarithmic scale in panel (d) further clarifies the behaviour: the model describes the horizontal shuttling data reasonably well, while in the vertical shuttling regime the agreement becomes progressively worse. In particular, at $\xi_{RF} = 0$ the model predicts vanishing heating, as expected for a single–drive configuration. This behaviour is consistent with the theoretical result from Subsection 3.2.2, which shows that phase noise does not contribute to heating when only one RF source is active. For $\xi_{RF} < 0$, the predicted ion heating increases, but remains below the measured values. Overall, the description along y shows better qualitative agreement with the measured data than the corresponding x–axis model.

These considerations suggest that the model captures the main features of the ion heating measurements performed during RF transport, but it remains incomplete. In-

deed, additional heating mechanisms not included in the present model may contribute to the measured ion heating and could explain the offset between the phase noise model and the measurements shown in panel (d) of Figure 5.11 for values of $\xi_{\mathrm{RF}} < 0.3$. In particular, noise on the DC electrodes or fluctuating patch potentials on the surface of the trap may be present in the experimental system. Although their individual contributions cannot be experimentally distinguished with the available measurements, they may nonetheless influence the observed dynamics and account for the remaining discrepancies at low ξ_{RF}. For higher values of ξ_{RF}, where one expects the heating due to the RF phase noise to be dominant over other sources of heating, the model indeed shows a better agreement with the data.

This concludes the comparison between the experimental data and the heating rate model developed in this work. Nevertheless, several additional considerations are required to fully understand the observed behaviour. First, the present analysis treats amplitude noise and phase noise independently, while in a real experiment both mechanisms may contribute at the same time. A more complete description would therefore combine the corresponding heating rates and fit a mixture of the two spectral noise densities. Second, the observation that the y-axis model provides a seemingly better description of the data in both noise scenarios suggests that the actual motional axes may not be perfectly aligned with the nominal trap axes. Simulations including mode coupling or small rotations of the radial modes could help determine whether such effects are sufficient to reproduce the measured heating rates. Finally, improved measurements — for example, the estimation of ion heating on a target motional mode via sideband thermometry [44] — will be essential to identify the contributions of different noise sources and to rigorously validate the proposed models.

6

Summary and outlook

The aim of this thesis was to develop a framework describing the relationship between RF technical noise and ion motional heating in trapped-ion quantum computing architectures, and to validate this framework through comparison with experimental data. In Chapter 3, a theoretical model was introduced to describe how technical noise on the RF drive couples to the motion of trapped ions, giving rise to motional heating. In Chapter 4, this model was used to characterize a fundamental building block of the QCCD architecture, the X-junction. In Chapter 5, the model was employed to study RF transport, which represents a core primitive of the QSA architecture, and the resulting predictions were compared with experimentally measured heating rates.

Within the context of ion motional heating induced by noise on the RF drive, two complementary models have been developed: one describing the effects of RF amplitude noise and the other describing the effects of RF phase noise. In both cases, the noise was treated as white noise for a resonator-driven circuit. These models establish a direct connection between technical noise originating from the RF electrodes and an effective fluctuating force acting on the ion, which gives rise to motional heating quantified in terms of heating rate. The models have been developed for both single- and multi-drive configurations. Notably, in the single-drive case with phase noise, the model predicts vanishing heating, as the noise is cancelled within the pseudopotential approximation; thus, only relative phase fluctuations between different RF driving signals contribute to motional excitation. In all other cases, the predicted heating rates depend on two quantities: the spectral noise density of the RF circuit and the gradient of the pseudopotential, which is determined entirely by the RF electrode geometry. Motional heating can therefore be mitigated by addressing these two contributions independently. The noise amplitude can be reduced by using RF resonators with higher quality factors and narrower bandwidths, effectively filtering out noise resonant with the ion motion. At the same time, the pseudopotential gradient can be minimized by

optimizing the shapes of the trap electrodes, thereby lowering pseudopotential barriers and regions of high pseudopotential variation.

In Chapter 4, the X-junction in the QCCD architecture was identified as a representative case in which pseudopotential barriers, reduced radial confinement, and increased motional heating arise during ion transport. To mitigate these effects, the RF electrode geometry was optimized using two parametrizations: piecewise- and spline-based approaches. In both cases, the pseudopotential barrier was significantly reduced, reaching approximately one third of the unoptimized value for the piecewise parametrization and one tenth for the spline parametrization. Radial confinement was also improved, with the piecewise approach yielding more stable secular frequencies, reduced only over a smaller region around the trap center, while both parametrizations outperformed the non-optimized design. Both approaches also showed reduced predicted motional heating, in both spatial extent (by roughly 30%) and amplitude (by about 50%) compared to the unoptimized geometry. Overall, for the geometrical dimensions considered, the spline-based optimization demonstrated slightly better performance than the piecewise approach, although the difference remained modest.

The obtained results confirmed that optimization of the electrode geometry is a key strategy for mitigating the impact of RF technical noise on ion motional heating. A natural extension of this work is the fabrication of the optimized electrode designs, followed by their experimental characterization. Along these lines, the design of the X-junction was finalized with a study of the DC electrodes and the development of a full-trap layout ready for fabrication, as presented in Section A.2. Moreover, direct measurements of motional heating rates and transport performance in the optimized junctions will be essential to validate the theoretical predictions and to assess their robustness under realistic experimental conditions.

In Chapter 5, the developed model was tested against experimental measurements of ion heating rates during RF transport, within the context of the QSA architecture. These measurements did not follow the standard behaviour expected from surface-noise induced heating. Simulations showed that RF technical noise can generate large and strongly varying heating rates even at constant ion–surface distance, a behaviour that cannot be explained by surface noise alone. Comparison with experimental data showed partial agreement. In the case of amplitude noise, the model predicts heating predominantly along the transport direction, although this directional separation cannot be directly extracted from the measurements. For phase noise, the predicted heating along the y-axis showed better agreement with the experimental data than the other cases.

Overall, the comparison indicates that RF technical noise can plausibly account for

the anomalous heating rates observed during RF transport. However, the model predicted ion heating along the axis perpendicular to the trap surface, which should not be accessible with the measurements performed. This suggests that the correspondence between the nominal trap axes and the actual orientation of the motional modes may not be exact, indicating the presence of small asymmetries or mode rotations in the trapping potential, possibly caused by stray fields present on the trap surface. Future work should include the estimation of ion heating on each motional mode using sideband thermometry during RF transport, as well as the extensions of the model to incorporate mode rotations. In addition, the optimization of electrode geometries should be explored for the QSA architecture, with the aim of identifying trap designs that are both experimentally feasible and capable of minimizing motional heating arising from RF technical noise.

A

X-junction: supplementary analysis

This appendix provides a supplementary analysis of the X-junction trap introduced in Chapter 4. The first section presents the optimized electrode layouts for different ion–surface distances. The second section examines the DC transport waveforms required to maintain stable axial confinement during shuttling along the junction. Finally, the last section introduces the problem of wire routing and its implications for practical device realization.

A.1 X-junction optimization results for different ion–surface distances

As discussed in Section 4.2, the geometry and dimensions of the RF electrodes in the X-junction directly determine the ion–surface distance. In particular, once a target ion–surface distance is selected, the corresponding geometrical parameters can be obtained using Equation 4.2, namely the gap between the RF electrodes, w_G, and their width, w_{RF}. The resulting values for different ion–surface distances are reported in Table 4.1. Different ion–surface distances can be chosen depending on the desired trade-off between trap depth, motional heating rates, optical access, and fabrication constraints. For instance, larger distances generally reduce anomalous heating and relax fabrication tolerances, while smaller distances allow for stronger confinement and more compact trap geometries.

The choice of the target ion–surface distance during shuttling along the junction also influences the subsequent optimization process. Although the optimization is performed using the adimensional pseudopotential defined in Equation 4.12, the electrode dimensions remain constrained by the selected ion–surface distance. As a consequence, different target heights lead to different optimized trap geometries.

For completeness, Figure A.1 and Figure A.2 show the optimized electrode layouts obtained using the piecewise and spline parametrizations, respectively, for two different ion–surface distances. In both figures, panel (a) corresponds to $h = 75\,\mu\text{m}$, while panel (b) corresponds to $h = 50\,\mu\text{m}$.

For the piecewise optimization, only minor differences in the electrode geometries

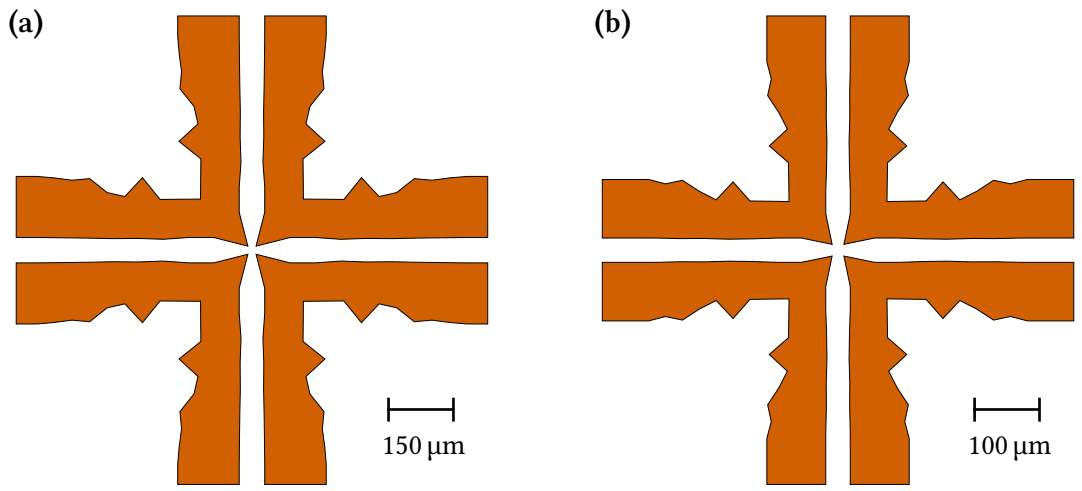

Figure A.1: Piecewise-optimized junction geometries. Panel (a) shows the result for $h = 75\,\mu\text{m}$, while panel (b) for $h = 50\,\mu\text{m}$.

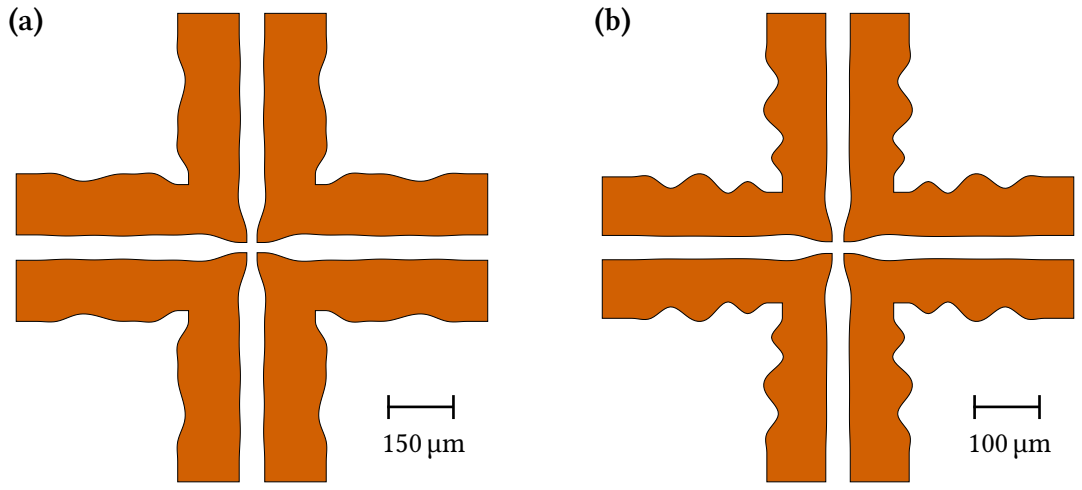

Figure A.2: Spline-optimized junction geometries. Panel (a) shows the result for $h = 75\,\mu\text{m}$, while panel (b) for $h = 50\,\mu\text{m}$.

are observed between the two cases. More pronounced differences are instead visible when compared to the geometry obtained for $h = 100\,\mu m$, shown in panel (b) of Figure 4.5. By contrast, the junctions optimized using the spline parametrization exhibit more significant variations in the electrode edge profiles. This behaviour is likely a consequence of the larger number of auxiliary points used to define the spline-based electrode shapes.

A pseudopotential analysis was also performed for all the considered geometries. The corresponding results are omitted here for brevity. In general, all optimized geometries exhibit a reduced pseudopotential barrier and less pronounced electric-field gradients with respect to the non-optimized configuration. For $h = 75\,\mu m$, the spline optimization tends to provide better performance, whereas for $h = 50\,\mu m$ the piecewise parametrization returns slightly improved results. In all cases, however, the achieved reduction of the pseudopotential barrier is smaller than that obtained for the configuration analyzed in Chapter 4 with $h = 100\,\mu m$.

A.2 DC transport waveforms

As anticipated in Section 4.1, the implementation of the QCCD architecture in planar surface traps relies on precise control of DC electrodes to generate the transport waveforms needed for shuttling ions through channels and junctions. Generally, the DC electrodes of a surface ion trap can be designed following two main approaches, distinguished by the scale of control they provide. In the *global control* approach, the electrodes are comparatively large, typically spanning several trapping sites. This configuration yields a smooth and globally defined potential landscape but offers limited flexibility for fine local tuning. Conversely, the *local control* approach employs electrodes whose dimensions are comparable to the ion–surface distance, allowing for precise manipulation of the potential in the vicinity of the trapped ion. Such finer control is essential for accurately shaping the axial confinement and enabling reliable ion shuttling operations. In all presented simulations, the trap has dimensions of $1500\,\mu m \times 1500\,\mu m$, with an electrode separation of $5\,\mu m$, for a target ion–surface $h = 100\,\mu m$.

Non-optimized junction

We begin the analysis by considering a non-optimized junction geometry featuring segmented inner DC electrodes and continuous outer DC electrodes, as shown in Figure A.3. The segmentation of the inner DC electrodes is intended to preserve the overall sym-

metry of the X-junction, thereby minimizing asymmetries in the trapping potential. In addition, segmenting the inner DC electrodes requires voltages approximately twenty times lower than those needed for segmenting the outer electrodes to achieve the same axial secular frequency [61, 80]. We consider the same electrode dimensions as those listed in Table 4.1 for a target ion–surface distance of $h = 100\,\mu\text{m}$. As shown in Figure 4.3, when shuttling through the non-optimized junction, the ion–surface distance is expected to vary between 100 µm to 170 µm.

The analysis aims to determine the DC voltage sets, referred to as *shim sets* (see Subsection 2.2.3), that independently control both the axial confinement frequency and the electric fields at the ion position, in order to compensate for micromotion. The shut-

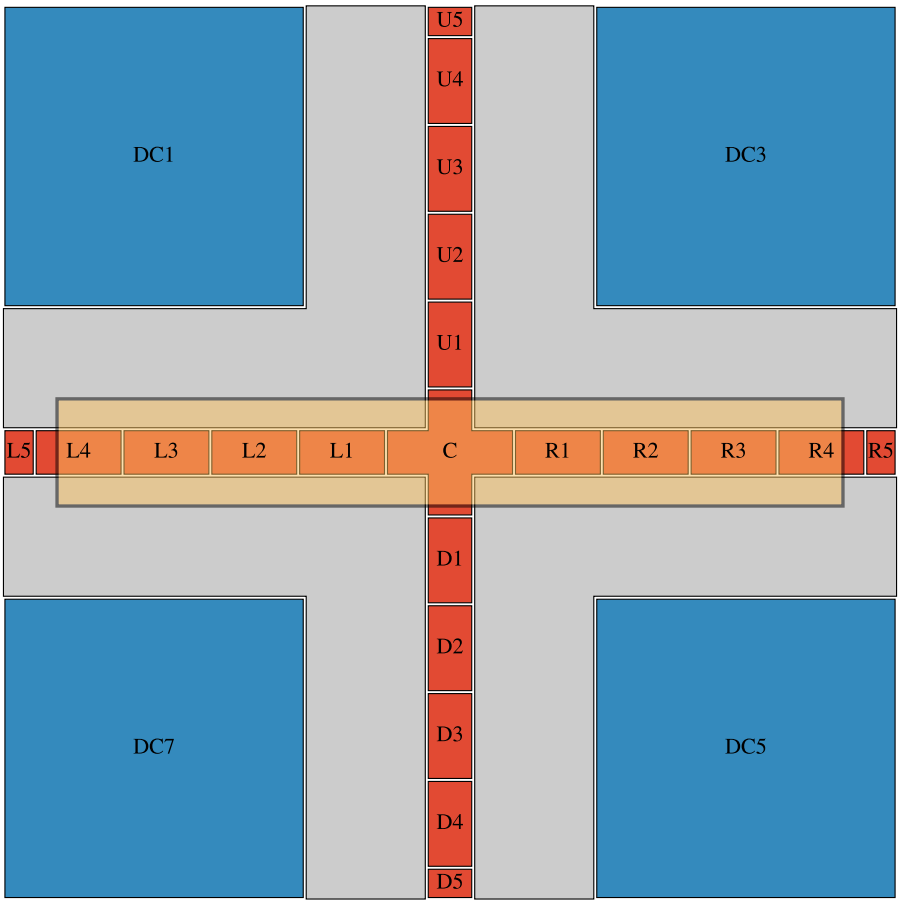

FIGURE A.3: Control electrodes of a junction design. The inner segmented electrodes (red) are divided in equal parts, except for the middle cross-like one. Each inner electrodes is indexed by the arm it belongs to ('U', 'D', 'L' or 'R') and its distance from junction center, e.g. U3.

tling region is highlighted in yellow in Figure A.3, and is chosen to avoid boundary effects that would not occur in a QCCD architecture, while the RF electrodes are kept at a fixed voltage and drive. At each axial position, the ion is assumed to reside at the local pseudopotential minimum. In computing the shim sets at a given position, all junction electrodes are included in the calculation, even those not immediately adjacent to the ion. The shim sets are calculated in order to achieve a total axial secular frequency of 1 MHz, which includes both the RF-induced axial confinement and the contribution from the DC voltages, and to produce electric fields of 1 V/mm at the ion position.

The results are shown in Figure A.4, where only the waveforms of electrodes reaching at least 1 V in absolute value are reported. These correspond to the electrodes belonging to the left and right arms of the junction, which are the ones involved in the considered transport path, while the electrodes of the upper and lower arms remain at comparatively low voltages and are therefore not displayed. The two shaded yellow regions again indicate the portions of the junction arm corresponding to the perpendicular RF

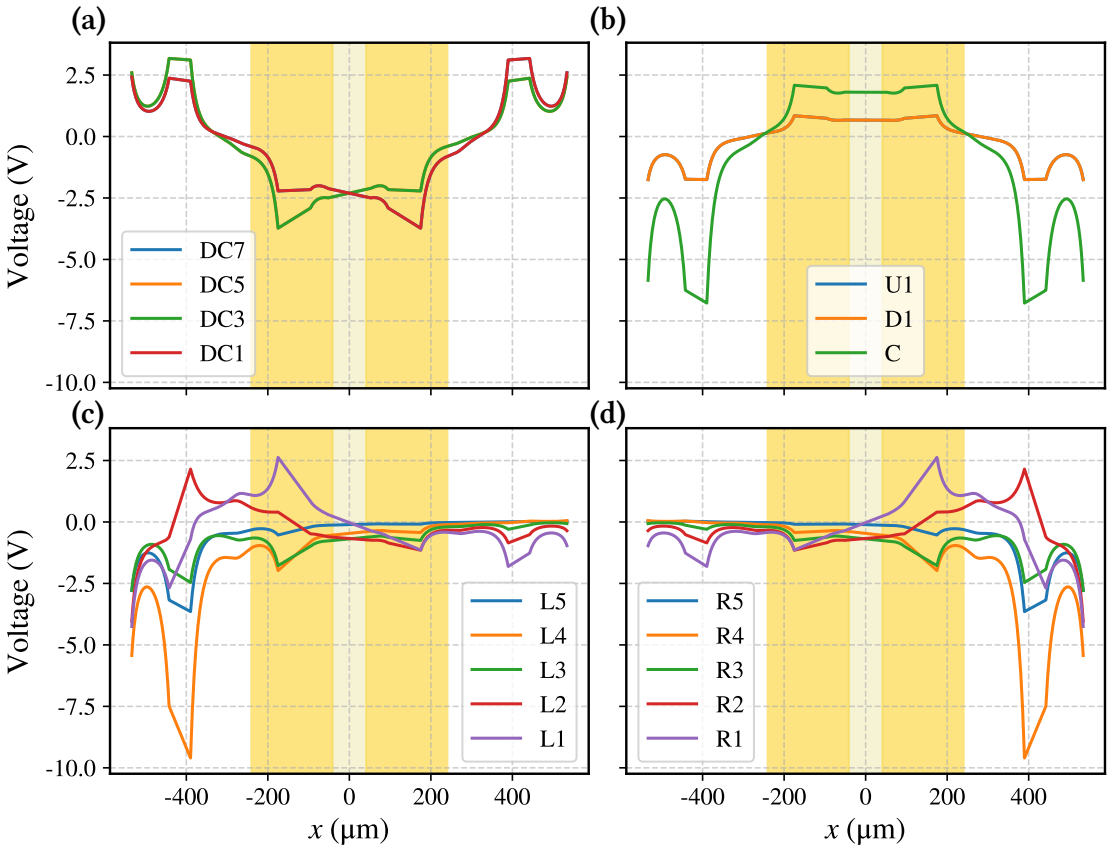

FIGURE A.4: Voltage waveforms for ion shuttling in the non-optimized junction. Panels show: (a) outer DC electrodes, (b) central and upper/lower arm electrodes, (c) left-arm electrodes, and (d) right-arm electrodes. Only waveforms exceeding ±1 V are displayed.

electrodes (darker area) and to the central opening of the junction (lighter area).

The first observation concerns the outer continuous electrodes, shown in panel (a). Due to the symmetry of the system, their waveforms are identical in pairs and exhibit positive peaks at $x = \pm 400\,\mu m$, corresponding to the small pseudopotential bumps seen in panel (a) of Figure 4.3. In the region of the perpendicular RF electrodes, the outer electrodes acquire negative voltages to compensate for the additional axial confinement produced by the RF fields. Panel (b) shows the central cross-shaped electrode, which behaves differently: it has negative peaks at $x = \pm 400\,\mu m$ and positive voltages near the trap center. Additionally, only the two upper- and lower-arm electrodes closest to the center, which share the same waveform, reach significant values, although still much smaller than the central electrode. The left- and right-arm electrodes, shown in panels (c) and (d) respectively, generally exhibit higher voltage amplitudes when the ion is nearby, with pronounced peaks at $x = \pm 400\,\mu m$. Finally, near the edges of the graph, boundary effects cause the voltages to rise slightly before the cutoff.

From these observations, we can draw the following conclusions:

- The transport through the junction can be achieved with feasible voltages (not exceeding $\pm 10\,V$), even when the junction is not optimized, while maintaining a fixed axial confinement of 1 MHz.
- The points at $x = \pm 400\,\mu m$ appear to be critical for ion transport, as the voltages reach their highest values. This suggests that these positions may be indicators of a threshold for the junction's impact on the shuttling parameters; however, this remains a hypothesis necessitating further examination.

Piecewise parametrized junction

The same analysis can be repeated for the optimized junction geometry obtained with the piecewise parametrization, introduced in Subsection 4.3.2. In this case, the inner DC electrodes are segmented as before, while the outer electrodes are also partially segmented in order to improve control during ion shuttling and enable micromotion compensation, as illustrated in Figure A.5. The highlighted region corresponds to the shuttling region considered for the analysis.

Following the same procedure adopted for the non-optimized case, we compute the DC transport waveforms required to shuttle the ion along the junction while maintaining a constant axial secular frequency of 1 MHz. The resulting waveforms are shown in Figure A.6. For clarity, only the waveforms of electrodes whose voltages exceed 1 V

in absolute value are shown. These correspond to the electrodes belonging to the left and right arms of the junction, which are the ones involved in the considered transport path, while the electrodes of the upper and lower arms remain at comparatively low voltages and are therefore not displayed. Furthermore, since the outer DC electrodes are identical in pairs, as previously mentioned, only a single representative waveform for each outer-electrode pair is shown in panel (a) to preserve readability. As in the non-optimized case, the shaded yellow regions mark the sections of the junction arm corresponding to the perpendicular RF electrodes (darker shading) and to the central opening of the junction (lighter shading).

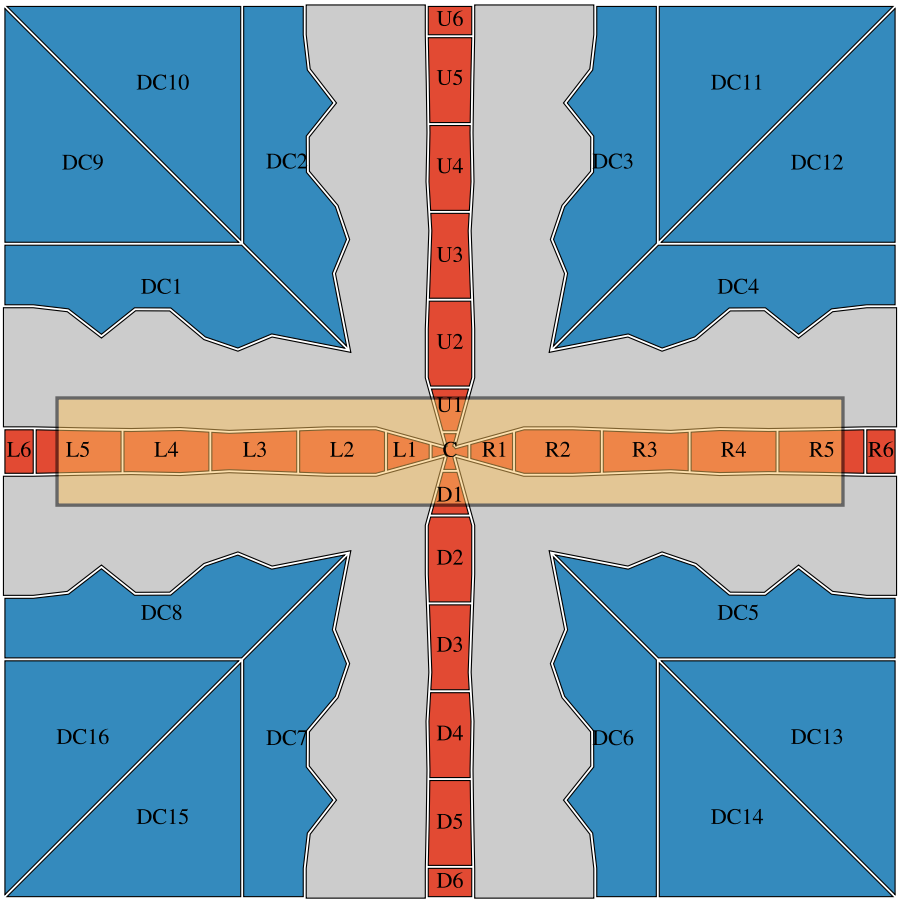

Figure A.5: Control electrodes of a piecewise-optimized junction geometry. The inner segmented electrodes (red) are divided in equal parts, except for the middle cross-like one. Each inner electrodes is indexed by the arm it belongs to ('U', 'D', 'L' or 'R') and its distance from junction center, e.g. U3. The outer electrodes are partially segmented to ensure better control.

The outer electrodes, shown in panel (a), reach higher voltages than in the non-optimized case and exhibit multiple peaks rather than a single maximum. This behaviour is likely a consequence of the inhomogeneous electrode edges, which are also responsible for the small oscillations observed in the axial secular frequency (see Figure 4.7). The most pronounced peaks are located around $x \approx \pm 500\,\mu\text{m}$. In the region of the perpendicular RF electrodes, the outer electrodes acquire negative voltages in order to compensate for the additional axial confinement produced by the RF fields. Panel (b) shows only the two upper- and lower-arm electrodes closest to the junction centre, which share the same waveform and are the only ones to reach significant voltage values. In this case, the central electrode (which is significantly smaller than in the non-optimized geometry, see Figure A.5) is not shown, as its waveform does not exceed 1 V. The left- and right-arm electrodes, shown in panels (c) and (d), respectively, generally display voltage amplitudes comparable to those shown in panel (a) when the ion

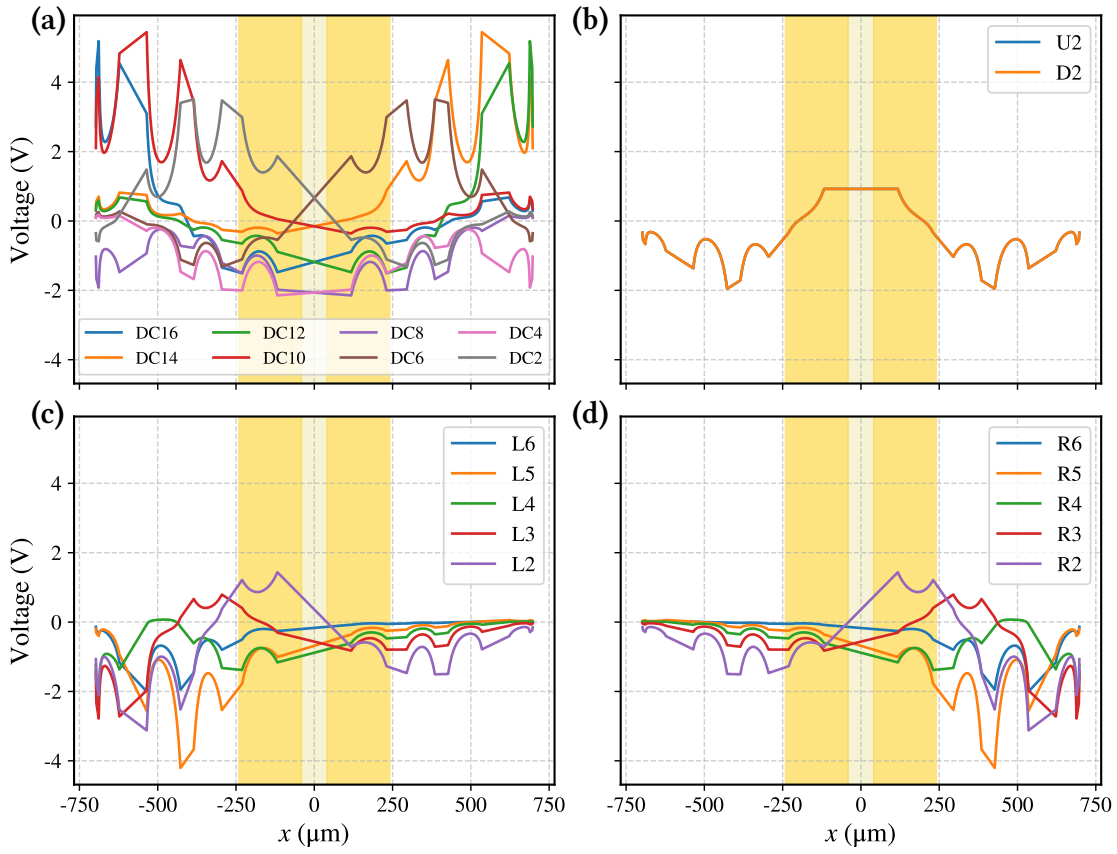

FIGURE A.6: Voltage waveforms for ion shuttling in the piecewise-optimized junction. Panels show: (a) outer DC electrodes, (b) central and upper/lower arm electrodes, (c) left-arm electrodes, and (d) right-arm electrodes. Only waveforms exceeding ±1 V are displayed.

is nearby, with maxima at $x \approx \pm 400\,\mu m$. Notably, the maximum voltages required are less than half of those needed in the non-optimized junction.

Spline parametrized junction

The same type of analysis can be performed for the optimized junction geometry obtained with the spline parametrization introduced in Subsection 4.3.3. As in the previous case, we adopt a DC electrode layout consisting of segmented inner electrodes and partially segmented outer electrodes, as shown in Figure A.7. The segmentation pattern is chosen to preserve the overall symmetry of the junction. The highlighted region

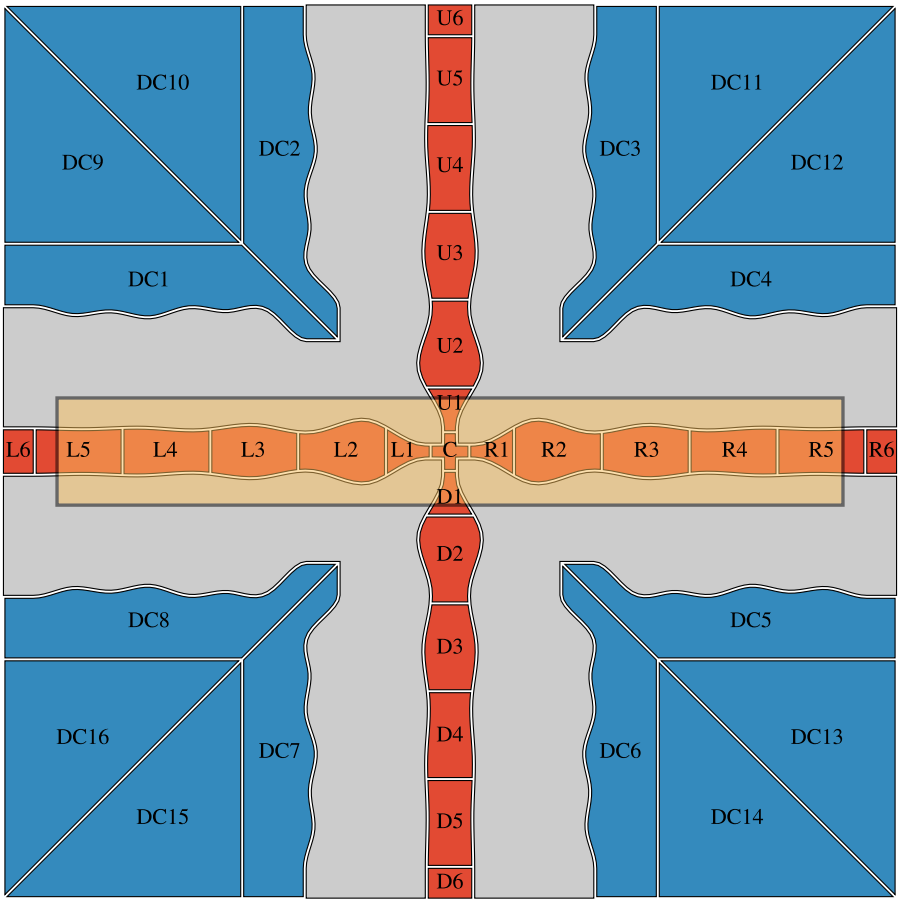

Figure A.7: Control electrodes of a spline-optimized junction geometry. The inner segmented electrodes (red) are divided in equal parts, except for the middle cross-like one. Each inner electrodes is indexed by the arm it belongs to ('U', 'D', 'L' or 'R') and its distance from junction center, e.g. U3. The outer electrodes are partially segmented to ensure better control.

corresponds to the shuttling region considered for the analysis.

Using the same procedure described earlier, we compute the DC transport waveforms required to shuttle the ion along the junction while keeping the axial secular frequency fixed at 1 MHz. The resulting waveforms are shown in Figure A.8, where only the waveforms of electrodes whose voltages exceed 1 V in absolute value are shown. The dominant contributions arise from the electrodes of the left and right arms, whereas the electrodes in the upper and lower arms remain near zero and are not reported. Owing to the symmetric pairing of the outer DC electrodes, only one representative waveform per pair is shown in panel (a). The shaded regions indicate the locations of the perpendicular RF electrodes (darker) and of the central junction opening (lighter).

The outer electrodes, shown in panel (a), reach higher voltages than in the other cases. Similarly to the piecewise parametrization, they exhibit multiple local maxima rather than a single peak, likely a consequence of the inhomogeneous electrode edges,

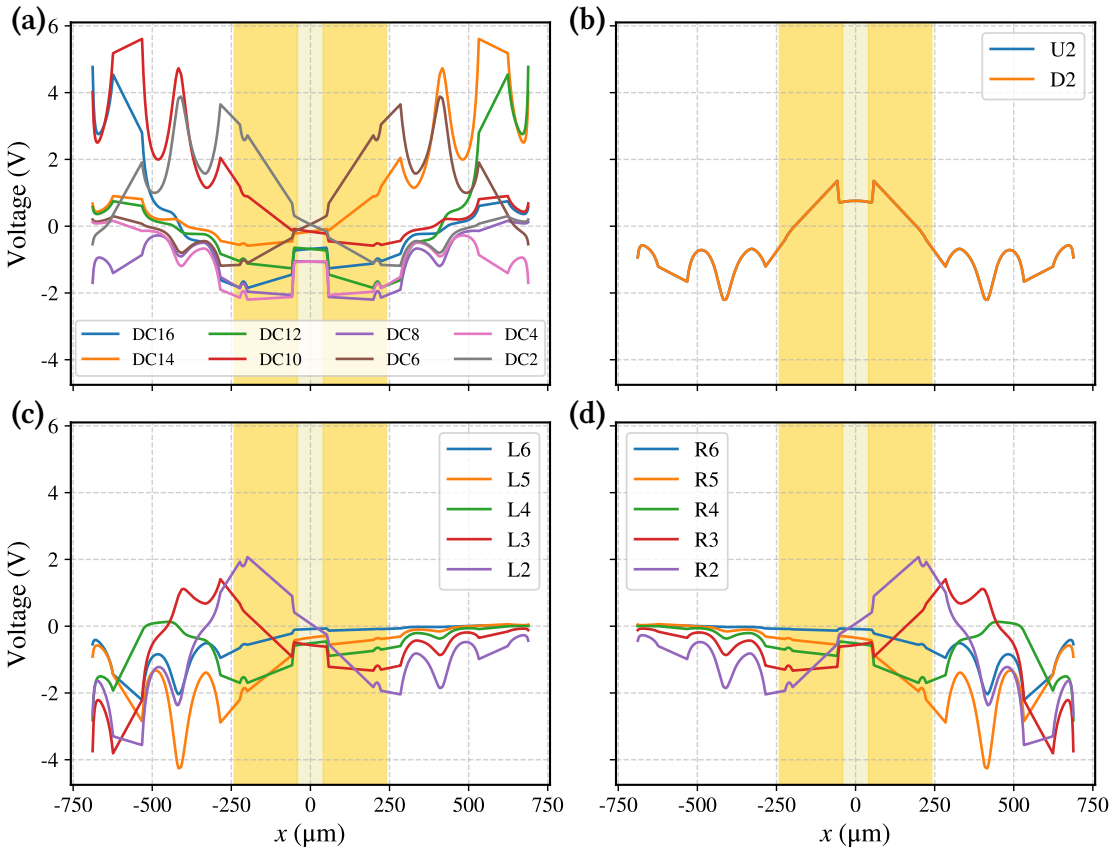

FIGURE A.8: Voltage waveforms for ion shuttling in the spline-optimized junction. Panels show: (a) outer DC electrodes, (b) central and upper/lower arm electrodes, (c) left-arm electrodes, and (d) right-arm electrodes. Only waveforms exceeding ±1 V are displayed.

which also give rise to the small oscillations observed in the axial secular frequency (see Figure 4.10). The most pronounced peaks occur around $x \approx \pm 500\,\mu\text{m}$. Panel (b) displays only the two upper- and lower-arm electrodes closest to the junction center, which share the same waveform and are the only ones to reach significant voltages. Also in this case, the central electrode is considerably smaller than in the non-optimized geometry (see Figure A.7), and its waveform does not exceed 1 V. The left- and right-arm electrodes, shown in panels (c) and (d), respectively, exhibit a behaviour very similar to that observed for the piecewise parametrization, with maxima at $x \approx \pm 400\,\mu\text{m}$.

Overall, these results indicate that the optimized junction, for both parametrizations, allows stable axial confinement to be achieved with lower voltages than in the non-optimized case, despite the waveforms being less smooth due to the spatial variations in the electrode geometry.

A.3 Wire routing

In a surface-electrode ion trap, each electrode must be electrically connected to its corresponding voltage source. For complex trap geometries, such as junctions, the wiring can become intricate, and it is important to design the connections in a way that avoids overlapping wires on the same substrate layer and minimizes overall complexity. We refer to this as the *routing problem*, that is to generate a layout for the routing of the trap electrodes that satisfies the requirements imposed by the desired connectivity and the trap manufacturing.

Pin zone designs

In modern surface-electrode traps, the connection of trap electrodes to external voltage sources is often achieved using *through-substrate vertical interconnect accesses* (TS-VIAs), i.e., vertical metal interconnects that pass through the substrate to link the electrodes on the trap surface to routing layers or to a PCB below. A straightforward design approach consists in placing a dedicated VIA directly beneath each electrode. While this guarantees a direct and simple electrical connection, it tightly couples the trap geometry to the substrate layout, making it difficult to fabricate the substrate independently of the trap design. This approach also becomes problematic for very small electrodes, where the required VIA can be comparable to or larger than the electrode itself and thus challenging to realize due to fabrication and dimensional constraints, limiting the feasibility and reliability of the connection.

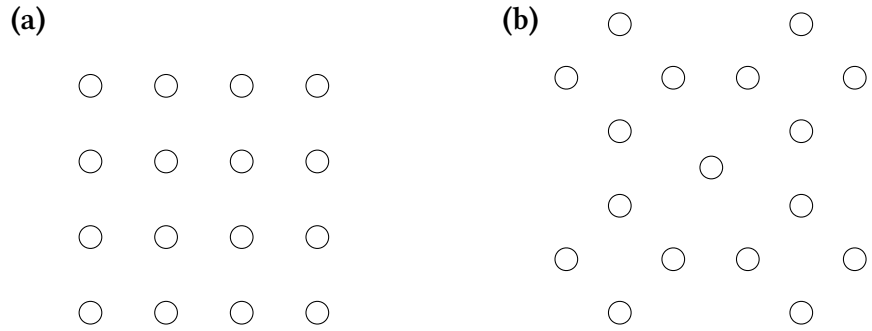

Figure A.9: Pin zone desings. (a) Simple 4×4 modular pin zone. (b) Modular pin-zone layout divided into four blocks, each consisting of a 2×2 array of VIAs rotated by 45°.

An alternative approach consists in grouping multiple TS-VIAs into dedicated connection areas, referred to as *modular pin zones*. In this scheme, several electrodes are routed towards shared connection regions, rather than being connected individually. This strategy decouples the substrate layout from the final trap geometry, enabling a more modular and flexible fabrication process. Several geometrical realizations of these modular pin zones are possible.

Assuming a square substrate, as in the X-junction case, pin zones can be positioned at the four corners of the chip. Within each zone, the VIAs can be arranged to minimise path overlap and reduce routing complexity, thereby improving wiring efficiency and layout regularity. Fabrication constraints must also be considered: VIAs cannot be placed arbitrarily close, and state-of-the-art microfabrication typically allows holes with a diameter of 30 µm and a minimum pitch (centre-to-centre distance) of 100 µm.

Two design examples for an ion–surface distance of 100 µm are shown in Figure A.9. Panel (a) depicts a simple 4×4 grid of pins. This layout is suboptimal because the inner pins can block access to the outer ones, complicating wire routing. Panel (b) shows a more practical arrangement, in which the pins are divided into four separate blocks, each rotated by 45°, leaving space for an additional central pin. Although this layout occupies slightly more space on the substrate, the staggered configuration reduces direct alignment between pins, simplifying the routing and making all VIAs accessible.

Routing optimization routine

Once the pin zone layout is fixed, the wire routing must be optimized. One effective approach is to employ a simulated annealing algorithm [99, 100]. Simulated annealing is a stochastic optimization method in which a candidate solution S is iteratively modified to explore the solution space. At each iteration, a new solution S' is generated in the

neighborhood of S, and the change in a cost function, $\Delta E = E(S') - E(S)$, is evaluated. The new solution is accepted with probability:

$$P = \begin{cases} 1, & \text{if } \Delta E \leq 0, \\ \exp(-\Delta E/T), & \text{if } \Delta E > 0, \end{cases} \tag{A.1}$$

where T is a control parameter called the *temperature*, gradually decreased according to a predefined annealing schedule. This strategy allows the algorithm to escape local minima early on and progressively converge towards a near-optimal solution as $T \to 0$.

In this context, the routing optimization is guided by two main objectives: eliminating wire crossings on a single substrate level (with the possibility of introducing additional levels to further reduce overlaps) and minimizing the overall wiring complexity to ensure manufacturability. The simulated annealing algorithm systematically explores possible routing configurations, iteratively refining the layout based on a cost function that penalizes crossings, excessive wire lengths, and unnecessarily convoluted paths.

An example of an optimized routing achieved with the simulated annealing approach is shown in Figure A.10, where the spline-optimized junction serves as a reference configuration. Two substrate layers, corresponding to the two colors, were employed to reduce overlapping of wires. Additionally, the corner electrodes located directly above the pin zones are connected without requiring additional routing, simplifying the overall layout.

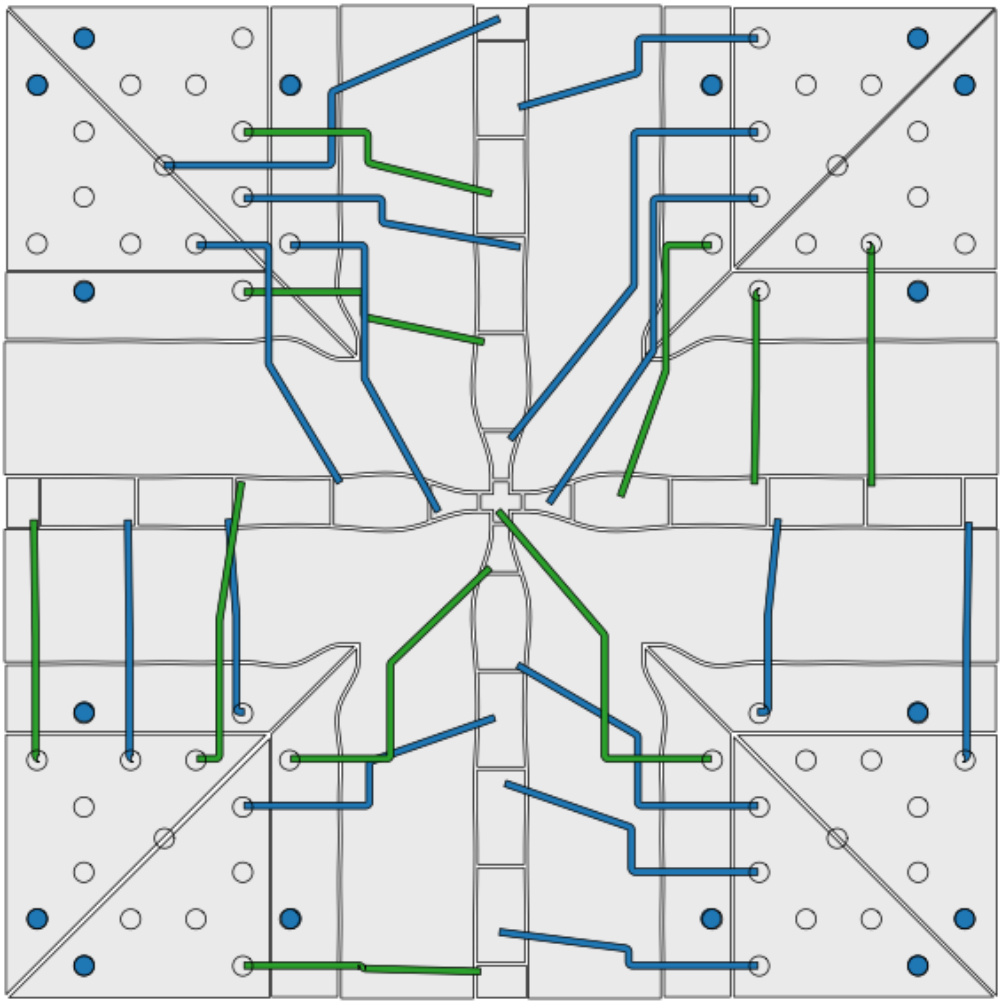

Figure A.10: Example of an optimized routing obtained with simulated annealing for the X-junction. The two colors represent two different routing layers.

B

Code development

The simulations presented in this thesis have been implemented in Python and organized in a publicly available release on GitHub [101], which provides a stable and citable snapshot of the code. The repository includes all source code, example notebooks, and scripts used to perform the analyses described in this work.

For the simulations, the electrode Python package was employed to compute the electric potentials and fields of the trap geometries. The gapless plane approximation, as discussed in Subsection 2.2.2, was used to model the surface-electrode traps, allowing efficient evaluation of the RF potentials and their gradients across the trap region. This combination enabled the numerical investigation of heating rates, RF transport, and junction optimization as presented throughout the thesis.

Bibliography

[1] A. M. Turing. "On Computable Numbers, with an Application to the Entscheidungsproblem". In: *Proceedings of the London Mathematical Society* 42.2 (Nov. 1936), pp. 230–265. ISSN: 0024-6115. DOI: 10.1112/plms/s2-42.1.230.

[2] R. P. Feynman. "Simulating physics with computers". In: *International Journal of Theoretical Physics* 21.6 (June 1982), pp. 467–488. ISSN: 1572-9575. DOI: 10.1007/BF02650179.

[3] P. Benioff. "Quantum Mechanical Models of Turing Machines That Dissipate No Energy". In: *Phys. Rev. Lett.* 48 (23 June 1982), pp. 1581–1585. DOI: 10.1103/PhysRevLett.48.1581.

[4] M. B. Plenio and V. Vitelli. "The physics of forgetting: Landauer's erasure principle and information theory". In: *Contemporary Physics* 42.1 (Jan. 2001), pp. 25–60. ISSN: 1366-5812. DOI: 10.1080/00107510010018916.

[5] P. W. Shor. "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer". In: *SIAM Journal on Computing* 26.5 (Oct. 1997), pp. 1484–1509. ISSN: 1095-7111. DOI: 10.1137/s0097539795293172.

[6] L. K. Grover. "A Fast Quantum Mechanical Algorithm for Database Search". In: *Proceedings of the Twenty-eighth Annual ACM Symposium on Theory of Computing (STOC '96)*. New York, NY, USA: ACM, 1996, pp. 212–219. ISBN: 0-89791-785-5. DOI: 10.1145/237814.237866.

[7] I. M. Georgescu, S. Ashhab, and F. Nori. "Quantum simulation". In: *Reviews of Modern Physics* 86.1 (Mar. 2014), pp. 153–185. ISSN: 1539-0756. DOI: 10.1103/revmodphys.86.153.

[8] B.-X. Wang et al. "Efficient quantum simulation of photosynthetic light harvesting". In: *npj Quantum Information* 4 (2018), p. 52. DOI: 10.1038/s41534-018-0102-2.

[9] M. Reiher et al. "Elucidating reaction mechanisms on quantum computers". In: *Proceedings of the National Academy of Sciences* 114.29 (July 2017), pp. 7555–7560. DOI: 10.1073/pnas.1619152114.

[10] M. H. Devoret and R. J. Schoelkopf. "Superconducting Circuits for Quantum Information: An Outlook". In: *Science* 339.6124 (2013), pp. 1169–1174. DOI: 10.1126/science.1231930.

[11] D. S. Weiss and M. Saffman. "Quantum computing with neutral atoms". In: *Physics Today* 70.7 (2017), pp. 44–50. DOI: 10.1063/PT.3.3626.

[12] X. Qiang et al. "Large-scale silicon quantum photonics implementing arbitrary two-qubit processing". In: *Nature Photonics* 12.9 (Aug. 2018), pp. 534–539. ISSN: 1749-4893. DOI: 10.1038/s41566-018-0236-y.

[13] L. M. K. Vandersypen and I. L. Chuang. "NMR techniques for quantum control and computation". In: *Rev. Mod. Phys.* 76 (4 Jan. 2005), pp. 1037–1069. DOI: 10.1103/RevModPhys.76.1037.

[14] A. Laucht et al. "Electrically controlling single-spin qubits in a continuous microwave field". In: *Science Advances* 1.3 (Apr. 2015). ISSN: 2375-2548. DOI: 10.1126/sciadv.1500022.

[15] D. J. Wineland et al. "Quantum information processing with trapped ions". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 361.1808 (July 2003), pp. 1349–1361. ISSN: 1364-503X. DOI: 10.1098/rsta.2003.1205.

[16] D. P. DiVincenzo. "The Physical Implementation of Quantum Computation". In: *Fortschritte der Physik* 48.9–11 (Sept. 2000), pp. 771–783. ISSN: 1521-3978. DOI: 10.1002/1521-3978(200009)48:9/11<771::aid-prop771>3.0.co;2-e.

[17] H. Haffner, C. Roos, and R Blatt. "Quantum computing with trapped ions". In: *Physics Reports* 469.4 (Dec. 2008), pp. 155–203. ISSN: 0370-1573. DOI: 10.1016/j.physrep.2008.09.003.

[18] W. Paul and H. Steinwedel. "Ein neues Massenspektrometer ohne Magnetfeld". In: *Zeitschrift für Naturforschung A: Physical Sciences* 8.7 (July 1953), pp. 448–450. DOI: 10.1515/zna-1953-0710.

[19] K. Kim et al. "Quantum simulation of the transverse Ising model with trapped ions". In: *New Journal of Physics* 13 (Oct. 2011), p. 105003. DOI: 10.1088/1367-2630/13/10/105003.

[20] P. Schindler et al. "A quantum information processor with trapped ions". In: *New Journal of Physics* 15.12 (Dec. 2013), p. 123012. ISSN: 1367-2630. DOI: 10.1088/1367-2630/15/12/123012.

[21] I. Pogorelov. "Quantum error correction in a compact ion-trap quantum computer". PhD thesis. Universität Innsbruck, 2023. URL: https://quantumoptics.at/images/publications/dissertation/pogorelov_PhD.pdf.

[22] K. Wright et al. "Benchmarking an 11-qubit quantum computer". In: *Nature Communications* 10.1 (Nov. 2019). ISSN: 2041-1723. DOI: 10.1038/s41467-019-13534-2.

[23] T. P. Harty et al. "High-Fidelity Preparation, Gates, Memory, and Readout of a Trapped-Ion Quantum Bit". In: *Physical Review Letters* 113.22 (Nov. 2014). ISSN: 1079-7114. DOI: 10.1103/physrevlett.113.220501.

[24] C. J. Ballance et al. "High-Fidelity Quantum Logic Gates Using Trapped-Ion Hyperfine Qubits". In: *Physical Review Letters* 117.6 (Aug. 2016). ISSN: 1079-7114. DOI: 10.1103/physrevlett.117.060504.

[25] P. Jurcevic et al. "Direct Observation of Dynamical Quantum Phase Transitions in an Interacting Many-Body System". In: *Physical Review Letters* 119.8 (Aug. 2017). ISSN: 1079-7114. DOI: 10.1103/physrevlett.119.080501.

[26] E. A. Martinez et al. "Real-time dynamics of lattice gauge theories with a few-qubit quantum computer". In: *Nature* 534.7608 (June 2016), pp. 516–519. ISSN: 1476-4687. DOI: 10.1038/nature18318.

[27] D. J. Wineland et al. *Experimental Issues in Coherent Quantum-State Manipulation of Trapped Atomic Ions*. June 1998. DOI: https://doi.org/10.48550/arXiv.quant-ph/9710025.

[28] D. Leibfried et al. "Quantum dynamics of single trapped ions". In: *Reviews of Modern Physics* 75.1 (2003), pp. 281–324. DOI: 10.1103/RevModPhys.75.281.

[29] K. Mølmer and A. Sørensen. "Multiparticle Entanglement of Hot Trapped Ions". In: *Physical Review Letters* 82.9 (Mar. 1999), pp. 1835–1838. ISSN: 1079-7114. DOI: 10.1103/physrevlett.82.1835.

[30] A. Sørensen and K. Mølmer. "Quantum Computation with Ions in Thermal Motion". In: *Physical Review Letters* 82.9 (Mar. 1999), pp. 1971–1974. ISSN: 1079-7114. DOI: 10.1103/physrevlett.82.1971.

[31] J. I. Cirac and P. Zoller. "Quantum Computations with Cold Trapped Ions". In: *Physical Review Letters* 74.20 (May 1995), pp. 4091–4094. DOI: 10.1103/PhysRevLett.74.4091.

[32] M. C. Smith et al. "Single-Qubit Gates with Errors at the 10^{-7} Level". In: *Physical Review Letters* 134.23 (June 2025), p. 230601. DOI: 10.1103/42w2-6ccy.

[33] A. C. Hughes et al. "Trapped-ion two-qubit gates with >99.99% fidelity without ground-state cooling". In: *arXiv preprint* arXiv:2510.17286v1 (2025). Submitted 20 Oct 2025. DOI: 10.48550/arXiv.2510.17286.

[34] D. Kielpinski, C. Monroe, and D. J. Wineland. "Architecture for a large-scale ion-trap quantum computer". In: *Nature* 417 (2002), pp. 709–711. DOI: 10.1038/nature00784.

[35] J. P. Home et al. "Complete Methods Set for Scalable Ion Trap Quantum Information Processing". In: *Science* 325.5945 (Sept. 2009), pp. 1227–1230. ISSN: 1095-9203. DOI: 10.1126/science.1177077.

[36] S. Seidelin et al. "Microfabricated Surface-Electrode Ion Trap for Scalable Quantum Information Processing". In: *Physical Review Letters* 96.25 (June 2006). ISSN: 1079-7114. DOI: 10.1103/physrevlett.96.253003.

[37] J. Chiaverini et al. "Surface-electrode architecture for ion-trap quantum information processing". In: *Quantum Information and Computation* 5.6 (2005), pp. 419–439. DOI: 10.48550/arXiv.quant-ph/0501147.

[38] D. Stick et al. "Ion trap in a semiconductor chip". In: *Nature Physics* 2 (2006), pp. 36–39. DOI: 10.1038/nphys171.

[39] D. J. Wineland and D. Leibfried. "Quantum information processing and metrology with trapped ions". In: *Laser Physics Letters* 8.3 (Jan. 2011), p. 175. DOI: 10.1002/lapl.201010125.

[40] Q. A. Turchette et al. "Heating of trapped ions from the quantum ground state". In: *Phys. Rev. A* 61 (6 May 2000), p. 063418. DOI: 10.1103/PhysRevA.61.063418.

[41] R. B. Blakestad et al. "High-Fidelity Transport of Trapped-Ion Qubits through an X-Junction Trap Array". In: *Phys. Rev. Lett.* 102 (15 Apr. 2009), p. 153002. DOI: 10.1103/PhysRevLett.102.153002.

[42] L. Deslauriers et al. "Scaling and Suppression of Anomalous Heating in Ion Traps". In: *Phys. Rev. Lett.* 97 (10 Sept. 2006), p. 103007. DOI: 10.1103/PhysRevLett.97.103007.

[43] D. A. Hite et al. "100-Fold Reduction of Electric-Field Noise in an Ion Trap Cleaned with In Situ Argon-Ion-Beam Bombardment". In: *Phys. Rev. Lett.* 109 (10 Sept. 2012), p. 103001. DOI: 10.1103/PhysRevLett.109.103001.

[44] M. Brownnutt et al. "Ion-trap measurements of electric-field noise near surfaces". In: *Rev. Mod. Phys.* 87 (4 Dec. 2015), pp. 1419–1482. DOI: 10.1103/RevModPhys.87.1419.

[45] M. Valentini et al. "Demonstration of Two-Dimensional Connectivity for a Scalable Error-Corrected Ion-Trap Quantum Processor Architecture". In: *Physical Review X* 15.4 (Nov. 2025). ISSN: 2160-3308. DOI: 10.1103/b9s1-6r44.

[46] S. Earnshaw. "On the nature of the molecular forces which regulate the constitution of the luminiferous ether". In: *Transactions of the Cambridge Philosophical Society* 7 (1842), pp. 97–112. URL: https://mit.edu/~kardar/research/seminars/Casimir2010/pdf/EarnshawPaper.pdf.

[47] M. W. van Mourik. "Control of multi-species trapped ions in a scalable quantum computer". PhD thesis. Universität Innsbruck, 2023. URL: https://quantumoptics.at/images/publications/dissertation/ThesisMvM_Web.pdf.

[48] E. Mathieu. "Mémoire sur le mouvement vibratoire d'une membrane de forme elliptique". In: *Journal de Mathématiques Pures et Appliquées, 2^e série* 13 (1868). Translation available via arXiv:2103.02730, pp. 137–203. DOI: 10.48550/arXiv.2103.02730.

[49] P. K. Ghosh. *Ion Traps*. Vol. 90. International Series of Monographs on Physics. Oxford, UK: Oxford University Press, 1995. ISBN: 9780198539957. DOI: 10.1093/oso/9780198539957.001.0001.

[50] F. A. Shaikh and A. Ozakin. "Stability analysis of ion motion in asymmetric planar ion traps". In: *Journal of Applied Physics* 112.7 (Oct. 2012). ISSN: 1089-7550. DOI: 10.1063/1.4752404.

[51] J. Heinrich. "A Be+ Ion Trap for H2+ Spectroscopy". PhD thesis. Sorbonne Université, Apr. 2018. URL: https://theses.hal.science/tel-02049379v1.

[52] F. G. Major, V. N. Gheorghe, and G. Werth. *Charged Particle Traps: Physics and Techniques of Charged Particle Field Confinement*. Vol. 37. Springer Series on Atomic, Optical, and Plasma Physics. Springer Berlin Heidelberg, 2005. ISBN: 978-3-540-22043-5. DOI: 10.1007/b137836.

[53] V. Roberdel et al. "Phase-space study of surface-electrode Paul traps: Integrable, chaotic, and mixed motions". In: *Physical Review A* 97.5 (2018), p. 053419. DOI: 10.1103/PhysRevA.97.053419.

[54] S. Hong et al. "Guidelines for Designing Surface Ion Traps Using the Boundary Element Method". In: *Sensors* 16.5 (2016), p. 616. DOI: 10.3390/s16050616.

[55] N. V. Konenkov, M. Sudakov, and D. J. Douglas. "Matrix methods for the calculation of stability diagrams in quadrupole mass spectrometry". In: *Journal of the American Society for Mass Spectrometry* 13.6 (2002), pp. 597–613. DOI: 10.1016/S1044-0305(02)00365-3.

[56] H. C. Nägerl et al. "Laser addressing of individual ions in a linear ion trap". In: *Physical Review A* 60.1 (1999), pp. 145–148. DOI: 10.1103/PhysRevA.60.145.

[57] T. F. Gallagher. *Rydberg Atoms*. Cambridge Monographs on Atomic, Molecular and Chemical Physics. Cambridge, UK: Cambridge University Press, 1994. ISBN: 9780521385312, 0521385318. DOI: 10.1017/CBO9780511524530.

[58] M. D. Hughes et al. "Microfabricated ion traps". In: *Contemporary Physics* 52.6 (Nov. 2011), pp. 505–529. ISSN: 1366-5812. DOI: 10.1080/00107514.2011.601918.

[59] M. A. Rowe et al. "Transport of Quantum States and Separation of Ions in a Dual RF Ion Trap". In: *Quantum Information and Computation* 2.4 (2002), pp. 257–271. DOI: 10.26421/QIC2.4-1.

[60] J. M. Amini et al. "Toward scalable ion traps for quantum information processing". In: *New Journal of Physics* 12.3, 033031 (Mar. 2010), p. 033031. DOI: 10.1088/1367-2630/12/3/033031.

[61] J. H. Wesenberg. "Electrostatics of surface-electrode ion traps". In: *Phys. Rev. A* 78 (6 Dec. 2008), p. 063410. DOI: 10.1103/PhysRevA.78.063410.

[62] M. G. House. "Analytic model for electrostatic fields in surface-electrode ion traps". In: *Phys. Rev. A* 78 (3 Sept. 2008), p. 033402. DOI: 10.1103/PhysRevA.78.033402.

[63] M. H. Oliveira and J. A. Miranda. "Biot-Savart-like law in electrostatics". In: *European Journal of Physics* 22.1 (Jan. 2001), pp. 31–38. ISSN: 1361-6404. DOI: 10.1088/0143-0807/22/1/304.

[64] R. Schmied. "Electrostatics of gapped and finite surface electrodes". In: *New Journal of Physics* 12.2 (Feb. 2010), p. 023038. ISSN: 1367-2630. DOI: 10.1088/1367-2630/12/2/023038.

[65] *electrode: Planar Electrode Field Solver*. https://pypi.org/project/electrode/. Version 0.1.7, accessed: 2025-03 – 2025-11.

[66] D. J. Berkeland et al. "Minimization of Ion Micromotion in a Paul Trap". In: *Journal of Applied Physics* 83.10 (May 1998), pp. 5025–5033. ISSN: 0021-8979. DOI: 10.1063/1.367318.

[67] R. Blatt and D. J. Wineland. "Entangled states of trapped atomic ions". In: *Nature* 453.7198 (2008), pp. 1008–1015. DOI: 10.1038/nature07125.

[68] C. L. Lawson and R. J. Hanson. *Solving Least Squares Problems*. Revised. Society for Industrial and Applied Mathematics, 1995. ISBN: 978-0-89871-356-5. DOI: 10.1137/1.9781611971217.

[69] R. B. Blakestad et al. "Near-ground-state transport of trapped-ion qubits through a multidimensional array". In: *Phys. Rev. A* 84 (3 Sept. 2011), p. 032314. DOI: 10.1103/PhysRevA.84.032314.

[70] S. Kogan. *Electronic Noise and Fluctuations in Solids*. Cambridge, UK: Cambridge University Press, 1996. ISBN: 978-0-521-46034-7. DOI: 10.1017/CBO9780511551666.

[71] N. Wiener. "Generalized Harmonic Analysis". In: *Acta Mathematica* 55.1 (1930), pp. 117–258. DOI: 10.1007/BF02546511.

[72] A. Khintchine. "Korrelationstheorie der stationären stochastischen Prozesse". In: *Mathematische Annalen* 109.1 (1934), pp. 604–615. DOI: 10.1007/BF01449156.

[73] R. Maiwald et al. "Stylus ion trap for enhanced access and sensing". In: *Nature Physics* 5.8 (2009). Published 28 June 2009, pp. 551–554. DOI: 10.1038/nphys1311.

[74] M. W. van Mourik et al. "rf-induced heating dynamics of noncrystallized trapped ions". In: *Phys. Rev. A* 105 (3 Mar. 2022), p. 033101. DOI: 10.1103/PhysRevA.105.033101.

[75] D. Trypogeorgos and C. J. Foot. "Cotrapping different species in ion traps using multiple radio frequencies". In: *Phys. Rev. A* 94 (2 Aug. 2016), p. 023609. DOI: 10.1103/PhysRevA.94.023609.

[76] C. J. Foot et al. "Two-frequency operation of a Paul trap to optimise confinement of two species of ions". In: *International Journal of Mass Spectrometry* 430 (2018), pp. 117–125. ISSN: 1387-3806. DOI: https://doi.org/10.1016/j.ijms.2018.05.007.

[77] W. K. Hensinger et al. "T-junction ion trap array for two-dimensional ion shuttling, storage, and manipulation". In: *Applied Physics Letters* 88.3 (Jan. 2006). ISSN: 1077-3118. DOI: 10.1063/1.2164910.

[78] D. L. Moehring et al. "Design, Fabrication and Experimental Demonstration of Junction Surface Ion Traps". In: *New Journal of Physics* 13.7 (2011), p. 075018. DOI: 10.1088/1367-2630/13/7/075018.

[79] A. Mokhberi, R. Schmied, and S. Willitsch. "Optimised surface-electrode ion-trap junctions for experiments with cold molecular ions". In: *New Journal of Physics* 19.4 (Apr. 2017), p. 043023. ISSN: 1367-2630. DOI: 10.1088/1367-2630/aa6918.

[80] C. Zhang, K. K. Mehta, and J. P. Home. "Optimization and implementation of a surface-electrode ion trap junction". In: *New Journal of Physics* 24.7 (July 2022), p. 073030. ISSN: 1367-2630. DOI: 10.1088/1367-2630/ac7db6.

[81] R. D. Delaney et al. "Scalable Multispecies Ion Transport in a Grid-Based Surface-Electrode Trap". In: *Phys. Rev. X* 14 (4 Nov. 2024), p. 041028. DOI: 10.1103/PhysRevX.14.041028.

[82] J. H. Wesenberg. "Ideal intersections for radio-frequency trap networks". In: *Phys. Rev. A* 79 (1 Jan. 2009), p. 013416. DOI: 10.1103/PhysRevA.79.013416.

[83] N. Konenkov et al. "Linear quadrupoles with added hexapole fields". In: *Journal of the American Society for Mass Spectrometry* 17.8 (2006), pp. 1063–1073. DOI: 10.1016/j.jasms.2006.03.013.

[84] K. Wright et al. "Reliable transport through a microfabricatedX-junction surface-electrode ion trap". In: *New Journal of Physics* 15.3 (Mar. 2013), p. 033004. ISSN: 1367-2630. DOI: 10.1088/1367-2630/15/3/033004.

[85] D. Gerlich. "Guided Ion Beams, RF Ion Traps, and Merged Beams: State-Specific Ion–Molecule Reactions at meV Energies". In: *AIP Conference Proceedings* 295 (1993), pp. 607–622. DOI: 10.1063/1.45274.

[86] J. D. Jackson. *Classical Electrodynamics*. 3rd. New York: Wiley, 1999. ISBN: 978-0471309321.

[87] J. M. Wilson et al. "In situ detection of RF breakdown on microfabricated surface ion traps". In: *Journal of Applied Physics* 131.13 (2022), p. 134401. DOI: 10.1063/5.0082740.

[88] Y Liu et al. "Cooperative engineering the multiple radio-frequency fields to reduce the X-junction barrier for ion trap chips". In: *Chip* 3.1 (2024), p. 100078. ISSN: 2709-4723. DOI: https://doi.org/10.1016/j.chip.2023.100078.

[89] R. T. Marler and J. S. Arora. "Function-transformation methods for multi-objective optimization". In: *Engineering Optimization* 37.6 (2005), pp. 551–570. DOI: 10.1080/03052150500114289.

[90] D. Hucul et al. *On the Transport of Atomic Ions in Linear and Multidimensional Ion Trap Arrays*. 2008. DOI: 10.48550/arXiv.quant-ph/0702175.

[91] A. Walther et al. "Controlling Fast Transport of Cold Trapped Ions". In: *Physical Review Letters* 109.8 (Aug. 2012). ISSN: 1079-7114. DOI: 10.1103/physrevlett.109.080501.

[92] H. Kaufmann et al. "Fast ion swapping for quantum-information processing". In: *Phys. Rev. A* 95 (5 May 2017), p. 052319. DOI: 10.1103/PhysRevA.95.052319.

[93] U. Utako Tanaka et al. "Design of a surface electrode trap for parallel ion strings". In: *J.Phys.B: At.Mol.Opt.Phys.* 47.3 (2014), p. 035301. DOI: 10.1088/0953-4075/47/3/035301.

[94] U. Tanaka et al. "Creation of double-well potentials in a surface-electrode trap towards a nanofriction model emulator". In: *Quantum Sci. Technol.* 6.024010 (2021). DOI: 10.1088/2058-9565/abe51a.

[95] P. C. Holz. "Towards ion-lattice quantum processors with surface trap arrays". PhD thesis. University of Innsbruck, Nov. 2019. URL: https://quantumoptics.at/images/publications/dissertation/Philip_Holz_PhD-thesis_server.pdf.

[96] R. J. Epstein et al. "Simplified motional heating rate measurements of trapped ions". In: *Phys. Rev. A* 76 (3 Sept. 2007), p. 033411. DOI: 10.1103/PhysRevA.76.033411.

[97] R. McConnell et al. "Reduction of trapped-ion anomalous heating byin situsurface plasma cleaning". In: *Physical Review A* 92.2 (Aug. 2015). ISSN: 1094-1622. DOI: 10.1103/physreva.92.020302.

[98] M. Niedermayr. "Cryogenic surface ion traps". PhD thesis. Universität Innsbruck, 2015. URL: https://quantumoptics.at/images/publications/dissertation/niedermayr_diss.pdf.

[99] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. "Optimization by Simulated Annealing". In: *Science* 220.4598 (1983), pp. 671–680. DOI: 10.1126/science.220.4598.671.

[100] E. H. L. Aarts and P. J. M. van Laarhoven. "Simulated annealing: An introduction". In: *Statistica Neerlandica* 43.1 (1989), pp. 31–52. DOI: 10.1111/j.1467-9574.1989.tb01245.x.

[101] Lorenzo Calandra Buonaura. *RF technical noise analysis in trapped-ion quantum computing architectures*. https://doi.org/10.5281/zenodo.17753565. GitHub repository, v1.0.0 release. 2025.

Acknowledgments

As this long dissertation finally reaches its end, it's time to acknowledge the many people I've crossed paths with along the way, those who, willingly or unwillingly, became part of this journey. The list is far from short, and deciding on the "right" order seems almost as impossible as finishing this dissertation itself. Still, I'll do my best, hoping not to leave anyone out, or at least not on purpose.

I first want to thank my previous supervisor, Carmelo Mordini, who introduced me to the world of ion-trapping quantum architectures and set me on the path that eventually led to this project. After his departure from the university, Professor Pietro Silvi kindly stepped in as my new supervisor, taking on the role at very short notice. I am grateful to him for his support and availability and for putting me in contact with the University of Innsbruck, where I carried out the internship that gave rise to this work.

I am especially thankful to all the people I met while living in the capital of the Alps. The five months I spent there profoundly shaped my studies and will likely influence my future as well. First and foremost, I would like to thank Philipp Schindler, my external supervisor, who welcomed me into the cryolab team in Innsbruck and followed my work closely. I deeply admire his knowledge and expertise, and I am particularly grateful for the time he always found to answer my questions and engage in discussions, no matter how busy he was. I also thank Thomas Monz, who helped me better understand both my project and the bigger picture, and who never missed an opportunity to poke me or to joke about Italians. I am grateful to Philip Holz and Martin van Mourik as well, for their time and availability in helping me resolve my doubts, and for the images I was able to use from their theses. My thanks also go to the rest of the cryolab team, whose warmth and support created an environment in which it was truly a pleasure to work. I would also like to thank Tommy, and my flatmates Mike and Reza, who made my time in Innsbruck truly special. Thank you for the time spent climbing, I am already looking forward to doing it again, for the long nights talking about anything and everything, and, of course, for the many pizzas we shared together. The last spot is reserved for Marco Valentini, whom I would almost consider a co-author of this thesis. A thousand words would not be enough to fully express my gratitude. From the very first time I came for the lab tour almost a year ago, he welcomed me with his sarcastic and joking attitude, which, as I soon discovered, hides a heart of gold. He followed my journey in Innsbruck like a big brother, constantly pushing me beyond my limits, helping me work

through my doubts, and teaching me how to express myself instead of hiding behind my fears. Throughout this time, he has been a true point of reference for me, showing how deeply one can care about another person and their growth. I can only hope that he finds his own path in the future, and I wish him all the best — good luck, Marco!

I also want to thank all the people I met in Padua, with whom I shared early Thursday morning classes at 8:30, countless hours in "LabWIFI," and numerous "aperitivi" after exams, no matter the hour. A special thanks to Giovanni, Vittoria, Claudia, Jacopo, Alessio, Chiara, Cesare, and Riccardo for making those moments memorable and full of laughter. My former flatmates, friends, and colleagues also deserve my gratitude: Pit, Fede, and Andre. I could list countless reasons, but above all, I thank you for sharing this adventure with me, living together in a new city, meeting new people, and embracing new opportunities. Thank you for the many political debates and endless laughter, for the hours spent studying in the kitchen, and for the nights spent enjoying (many) beers together.

The people I left in Modena must also be thanked. Thank you to Luca, who always made an effort to meet whenever possible, and for the videocalls and messages that kept me company while I was away, you are a true friend. Thanks to the "Biretta" group for always keeping a spot open for me, even after I left the city. My theatre friends deserve a special mention as well. Each one for a different reason, but to list them all properly would take too long. Thank you, Cave, Morgan, Cewa, Sam, Raffa, Asca, Momo, Rosi, Mary, Tea, Gre, Ale, Ali, Anna, Mati, Sara and Elena, you are like a family. Even as we grow up (and some of us buy houses...) and scatter across the world, there will always be a red thread connecting us, and I am sure we will always know how to find each other again. Thank you for the holidays, the shows, and the nights spent at the "Goblet", you are an irreplaceable group that will always hold a place in my heart.

This whole journey would not have been possible without my family, who have always supported me in my studies and my choices. My mum, for the kindness and concern that define her, for her constant calls and messages, and for never letting me leave the house without some good food. My father, for all the practical help, especially when moving to Padua and Innsbruck, and for his quiet but steady support. My brother, for always being proud of me and for standing by me throughout these years. And finally, my two dogs, Bianca and Lucky, who never fail to remind me what pure enthusiasm looks like, always ready to play and cover me in kisses as if I had never left.

Last but not least, Lucia. You have been the biggest part of my life for more than six years, and you have seen me through every step of my academic journey, from high school, to my Bachelor's degree, and now this Master's. You are always by my side,

supporting me and showing me love, making everything a little easier and a lot more enjoyable. Distance does not define us; if anything, it makes our bond stronger. Every time I leave, I know that you will always be there for me, waiting for new adventures, or for yet another crazy idea of yours. Wherever the next chapter will take me, I can only hope that it keeps finding its way back to you, my λόγος.