UNIVERSITÉ SORBONNE PARIS NORD
INSTITUT GALILÉE
LABORATOIRE DE PHYSIQUE DES LASERS

Thèse présentée par

Joseph Seaward

Pour obtenir le grade de

Docteur de l'Université Sorbonne Paris Nord

Sujet :

Unidimensional gas with tunable interactions

Soutenue le 20 novembre 2020 devant le jury composé de :

Mme	**Patrizia VIGNOLO**	**Rapporteuse**
M.	**Carlos GARRIDO ALZAR**	**Rapporteur**
Mme	**Juliette BILLY**	**Examinatrice**
M.	**Sébastien CHÉNAIS**	**Examinateur**
M.	**Maxim OLSHANII**	**Examinateur**
Mme	**Hélène PERRIN**	**Directrice de thèse**
M.	**Aurélien PERRIN**	**Membre invité, co-encadrant**

Contents

Acknowledgements

I begin by thanking Maxim Olshanii, whose Mathematical Methods for Physicists class I more-or-less walked into off the street — he has stuck by me since then; telling me about this PhD position and even introducing me to the problem in the second part of this thesis. I would also like to thank Hélène Perrin for being so generous with her time, energy and home as well as applying the pressure making those demands of exactness that a serious student needs. And then there is Aurélien Perrin, whose kindness and availability should make him a legend among graduate students forever.

I would also like to thank the rest of my committee: Juliette Billy, for her kind and insightful questions, Sébastien Chénais for presiding over my defense, and in particular the two reporters, Patrizia Vignolo and Calos Garrido Alzar for their thorough reading of this thesis.

Not to be forgotten are those without whose help none of this would be possible, the administrative team: Maryse, Carole, and the skilled Solen.

Thanks go out to Thomas Badr and also an apology for always forgetting how funny he is and so missing his jokes.

Then there are the comrades-in-science, David Rey chief among them, who have made this work something far different from a chore: Ali, Karim, Louis, Biplab, Yanliang, Mohammed, the poet Guesswho, Mathieu, Kaci, and Christelle. Callum Duncan deserves particular mention for the sound advice and warm welcome he offered me from the store of these things laid up in his home country.

There are those scientists who outrank me, whose encouragement, kindness, advice and support have helped me more than they realize: Félix Werner and Jim Bales.

I have had some teachers, not directly involved with this work, who nevertheless I would not be doing it without: Steve Arnason and Bala Sundaram. Maxim told me that you put in a good word for me.

Thanks also to the people not sharing in the struggle of physics, French and otherwise, who made me feel welcome here: Kenzy, Victor, Miguel, Delphine, and Vanessa with her ant lions.

Then there are my own ones: of whom Lilly is the first, whose letters have been such a nourishing treat while I have been doing this work, followed by Molly and Oscar, Xa, and Ian.

Thanks S.C. for the sweater, and everything.

And Mom, of course.

General Introduction

This thesis treats two separate but related pieces of work dealing in different ways with one-dimensional systems. One part is experimental work, following on from the thesis of D. Ben Ali [1], to bring online an ultra-cold sodium gas experiment. This involved preparing mounting and testing our atom chip; designing imaging objectives for absorption imaging; and optimizing a magnetic transport based on chained coils modeled after the one introduced by Greiner et al. [2]. Vacuum problems with the experiment left me time for a small theoretical collaboration with M. Olshanii which is dealt with in the other part. This work begins formulating a way to study the ergodicity [3] of one-dimensional systems by defining a certain kind of counter over dynamical systems [4] and determining the growth of that counter. While they both relate to one dimensional systems, they arise out of fairly different scientific contexts so I will have to recount their lineages separately.

Background of the experimental work

The phenomenon of Bose–Einstein condensation — macroscopic occupation of the ground state due to the saturation of excited states at low temperature — was first described in two papers by A. Einstein in 1924-5 [5, 6]. He was looking at the consequence of a method of deriving Planck's law published by S. N. Bose in 1924 [7]. Bose's derivation relies on the quantum statistics that arise when exchanging particles does not give rise to a new state. This is a distinguishing property of the particles which bear his name, *bosons*, which have integer spin, may share a quantum state, and whose many-body wave function is symmetrical under particle exchanges. These stand in contrast to fermions which have half-integer spin, obey the Pauli exclusion principle [8], and whose many-body wave-function is anti-symmetric under exchange of particles.

Einstein described this condensate phase of matter only in principle; it would occur with particles which followed Bose's statistics at very low temperatures. However, he could not give a concrete example in Nature. In 1938 F. London proposed that a phase transition to a superfluid in liquid helium, known to occur at 2.19 K, *was* the phase transition predicted by Einstein [9], albeit distorted by interaction forces which were neglected in Einstein's description, his being the treatment of an ideal gas. It would be many decades until experimental techniques advanced enough to produce condensates from the gas phase and allow experimenters to manipulate and probe them.

Those experimental breakthroughs began in the 1980s with the development of the Zeeman

slower for slowing jets of atoms [10] and methods for cooling and trapping atoms [11–13]. These developments won S. Chu [14], C. Cohen-Tannoudji [15], and W. Phillips [16] the 1997 Nobel Prize in Physics. With these techniques (and one or two others [17,18]) experimentalists could trap and cool dilute gases of atoms to very low temperatures. This led to the first condensate from a gas in 1995 when the group of C. E. Wieman and E. Cornell formed a condensate of rubidium-87 atoms [19], followed shortly by W. Ketterle's group making one from sodium that same year [19]. These three were awarded the 2001 Nobel Prize in Physics for these accomplishments [20, 21]. Since then many atoms, including the formidable cesium atom [22], have been condensed as well as molecules [23], phonons [24] magnetons [25], photons [26], and exciton polaritons [27, 28]. Aside from studying basic physics, they have been used for applications as widely ranging as information processing [29] and simulating the expanding universe [30].

A large part of the appeal of Bose–Einstein condensates is that they are systems of macroscopic size with their quantum effects intact. In the case of condensates of alkali atoms, they offer a variety of ways for the experimenters to control and interact with them [31]. One important mechanism of control they make available is control over the strength of their interactions. This is typically done via the mechanism called Fano–Feshbach resonance [32]. This also has its roots in early twentieth century physics, going back to O. K. Rice explaining the broadening of certain spectral lines due to transitions that occur when the potential energies of different electron states are equal [33]. This was extended to the coupling of discrete, bound states to states where the atoms may separate by H. Feshbach [34] and U. Fano [35] in 1958 and 1961, respectively. A way such resonances might arise in quantum gases in magnetic fields dates back to 1976 and W. Stwalley [36], well before the creation of the first trapped condensates formed from dilute gases. Once the cooling and trapping techniques began to develop, the use of these as a method of control was pointed out by Tiesinga et al. in 1993 [37] and was experimentally demonstrated in 1995 by Moerdijk et al. [38] with sodium and by Courteille et al. [39] with rubidium. Feshbach resonance can change attractive interactions into repulsive ones, and vice-versa, and was key to making the condensate of cesium stable [22]. It has become a standard tool and remains an object of study [40]

But magnetic fields are not the only way to couple bound and free states. It can also be done with photons [41,42]. In 2010, Papoular et al. introduced a Feshbach resonance induced by microwaves [43]. This has yet to be experimentally confirmed and it is something we not only hope to demonstrate, but to use to control the atoms in our experiment. That is part of our motivation for choosing to make our final trap with an atom chip [44,45]. These are chips made from layers of gold, laid down in wires of various configurations, that produce magnetic fields that can trap and control the atoms. They can accomplish two important things for us: incorporate a coplanar microwave guide [46] to deliver the microwaves to tune the interaction strength; and produce very elongated magnetic traps that will confine the dynamics of the gas to one dimension [47].

What we are interested in testing are theoretical predictions about systems where motion or excitations can only occur along one direction. When bosons in one dimension have interactions strong enough that the interaction energy is sufficient to correlate the particles (excluding the possibility that two or more can be found at the same spot) well-known mean-field theories break down. However, in this strongly-interacting regime there are exact solutions such as the Lieb and Liniger model [48, 49] and its extension to finite temperature by Yang and Yang [50]. There has yet to be a direct, experimental test of these solutions — one which studies a single, one-dimensional gas of strongly-interacting bosons at uniform

density. Tests have been done with arrays of gases [51–53], which change either the longitudinal confinement or the depth of a lattice on the longitudinal axis. These are capable of probing the strongly-interacting regime after one averages over the array. Tests on chips have just barely reached the strongly interacting regime in an individual tube in the group of I. Bouchoule who used a current-modulation technique to bring the cloud very close the surface of her chip, producing very strong gradients, while overcoming the roughness of the wires that would interrupt the potential if the current through them was direct [54]. However there are limitations on this method [55] and they could not go very far into the strongly-interacting regime. One of the goals of this experiment will be to probe deep into the strongly interacting regime by tuning the interaction strength of the atoms with this novel Feshbach resonance, independent of the density, and make comparisons to the models as exact as possible.

The different regimes of the one-dimensional gas are distinguished by differences in their second-order correlation function [56]. This is one way to distinguish the strongly-interacting gas from the ideal gas from a "quasi-condensate," [57] all of which have interesting physical characteristics I will talk about later in this thesis. Because they are distinguished by correlations in fluctuations, many instances must be taken in order to build up meaningful statistics. Thus, minimizing the time between preparations of the gas is a worthwhile investment of experimental energy and is one of the reasons that our atoms undergo initial cooling in a chamber separate from our atom chip. That way we can cool one sample while we perform experiments on another. This means that the atoms must be transported from one chamber to the other and the optimization of that transport was a significant part of this thesis project.

Background for the theoretical project

The actual loading of the chip and preparation of a one-dimensional sample was significantly delayed by vacuum problems in the science chamber containing the chip. The Helicoflex^{TM} gaskets used for the windows did not sustain the bake-out process very well, but it was many months before we found that they were the point of failure. This left me with some time to devote to a theoretical and numerical project with my old adviser, M. Olshanii, related to ergodic theory.

Ergodicity is a question that goes back to L. Boltzmann [58]. He says himself that he considers the probability of a system being found in a certain state to be the average amount of time it spends in that state, while J. C. Maxwell (whose work he is commenting on) considers it the portion of an ensemble of equivalent systems in all possible initial states.[1] To say that a system is ergodic is, then, to say that these views are equivalent. Boltzman knew he could prove this if he could show that in infinite time, the phase trajectory of the gas visited every point in its phase space. It was pointed out to him that it is impossible for a one-dimensional trajectory to fill a higher dimensional space completely and it was later proved that this specific notion of ergodicity is impossible under Newton's laws [60, 61]. But the notion of "quasi-ergodicity" (of a trajectory densely filling the phase space after infinite time [62]) was shown by J. von Neumann [63] and G. Birkhoff [64] to be sufficient for the convergence of the averages in infinite time. When considering the ergodicity of quantum systems this must be formulated slightly differently [65, §7.2] since the natural objects in quantum systems are operators and Hilbert spaces, instead of functions defined on phase spaces.

This is a key aspect in the study of dynamical systems, an exceedingly broad topic. It comprises systems which can be described by a phase space and an updating rule which maps

1. Paraphrased from [59, p.582] as quoted in [58].

each point in the phase space to the point which is "next" in time. As you may guess this encompasses a great many physical systems. A smaller and simpler subset of these systems are know as billiards [66], which can be thought of as a point moving at a constant velocity inside of a shape, elastically colliding with the sides. These can model one-dimensional systems by a coordinate transform which maps collisions with the sides to collisions between the one-dimensional particles. A billiard of dimension N can encode a system of $N-1$ particles and in certain cases [67] exact quantum solutions can be obtained for these systems [68] which, unlike the Lieb–Liniger solution, can contain a mixture of masses and interaction strengths (See [69] and references therein.).

The inspiration for the theoretical project presented here came from a paper by Wang et al. [70] where they count events in a billiard in a particular way. Over time this counter tends to stay near its initial value and hence seems to suggest that the system is not ergodic. It must be said that counters are not observables (they *count* events in the phase space, they are not *functions* of the points in phase space) and the fact that they do not behave ergodically is not proof that the underlying system is not ergodic. Still, arguments can be made from how the counters grow to how averages of the observables converge. We hope to develop a family of counters that can tell us by their growth whether a system or observable is ergodic. The work is in the early stages and what is presented here is a rule by which the kind of counter we are interested in can be expected to grow.

Outline of this thesis

Part I covers the experimental work in five chapters while Part II covers the theoretical and numerical work in two chapters, not counting a brief introduction.

Chapter 1 covers the theoretical basis for the experiment. I begin with the Bose–Einstein condensation of an ideal gas in three dimensions, add contact interactions, then reduce the dimension, and discuss some of the physics to be expected there. Most important are the different regions of the phase diagram of the one-dimensional gas that we want to explore, so I discuss the behavior of correlations in those regions. Then I outline the mechanism by which we will change the interactions and reach these different regions: the microwave-induced Feshbach resonance.

Chapter 2 gives an outline of the experimental setup, its current state and what will be done with it. It is meant to let the reader understand the rest of Part I without looking up the thesis of D. Ben Ali [1] which covers the design and construction in much more detail. It will make clear the motivations for our experimental choices and cover the relatively minor updates put in place since D. Ben Ali's graduation. This chapter will also cover what we intend to measure and how, since this has direct bearing on the objective lenses I designed.

Chapter 3 details the atom chip, fabricated for us by M. Stüwe in the group of J. Schmiedmayer. It covers how it forms the fields for magnetic trapping, cooling, and microwave dressing, as well as the design and construction of the chip and its support. The main result of this chapter is the outcome of the electrical tests performed on the chip; they are presented along with everything necessary to estimate what kind of fields and traps the chip can be expected to generate.

Chapter 4 describes the designs for imaging objectives I prepared for imaging the atoms in the science chamber. There is a low-magnification one, and a high-magnification one that can be converted to an even higher magnification for imaging the atoms still in the trap on the chip. To begin I present the basic optics needed to understand the compromises that had

to be made in designing the objectives. Then I go over how those compromises were made and what software simulations suggest will be the performance of the objectives. Performance tests are being carried out now so data are not yet available to check the simulations.

Chapter 5 presents the work done optimizing the magnetic transport. The theory of operation for these types of transports is reviewed and the loss mechanisms we must deal with are discussed. What we optimized specifically was the trajectory we used to move the atoms; the ones that make for the most relevant comparison are defined, with some others being included in Appendix D. Simulations were used to distinguish between loss mechanisms for the different profiles. I review them, and what they tell us about the experimental data. Details about the thermodynamics of the gas in the magnetic trap are also given here, as we want to transport the cloud while heating it as little as possible.

Chapter 6 is the first full chapter of the second part. It contains all the background and definitions needed to understand the work detailed therein, starting with what I will rely on from ergodic theory. I then deal with billiards, making links to one dimensional systems where I can. Then I talk about the dynamical system we used as a test case: irrational rotations. This involves talking about rotations generally, and how we approximate irrationals.

Chapter 7 gives the meat of the theoretical and numerical project. The case that inspired it is analyzed, and our numerical methods are discussed and compared. The counter we studied is defined, and we try to estimate its growth and also its range of validity. A rule for counters of this type is found that will hold for irrationals that can be well-approximated by rationals (a notion I define in Chapter 6). Difficulties and prospects for generalizing our results are dealt with here.

Finally I review my results and put them in their scientific contexts and discuss the future prospects for these projects.

Part I

Sodium experiment

Chapter 1

Introduction: a one-dimensional quantum gas with tunable interactions

Part I of this thesis is devoted to an experiment whose goal is the production of a single species quasi-condensate in a single (or pair of) strongly confining traps which reduce the dynamics to one dimension with controllable interaction strength. In this chapter I will lay out the theory necessary to define and discuss such a system, and give a picture of the physics we hope to explore. This theory, along with some sketches of the current state of the art, will make the motivations for our experimental choices clear, including our decision to work with sodium.

First I will give the basics I will need regarding quantum gases and condensation in three dimensions. I will work from the most general case of non-interacting bosons and then add in the elements we will need to discuss one-dimensional, contact-interacting systems and their phases. For my most basic facts, I will draw mostly from the textbooks of Landau and Lifshitz [71, 72]. For things more particular to cold atoms and the general structure of Section 1.1 I follow [73], but borrow from a lecture by H. Perrin [74]. Then I will deal with the specifics of those one-dimensional gases, the models, and the previous experiments we will be relying on for guidance as we explore the phase diagram. Then I will treat the specific tool we will be using to do that, the as-yet unobserved microwave-induced Fano–Feshbach resonance [43]. This is an extension of the more well-established procedure of using Fano–Feshbach resonances [33–35] [32, for a review] based on magnetic fields [38] and optical couplings [41] [42, for a review] to tune the interaction strength. This, along with laser [10, 11, 75] and evaporative [19] cooling techniques will allow us drive the gas into its different regimes.

For this chapter I rely on two excellent reviews: one by Bouchoule et al. [47], particularly useful for its review of the experimental work being done on atom chips; and one by Bloch et al. [76] with an extensive bibliography of seminal papers on many topics. Other reviews and textbooks will provide the structure for specific sections and I will cite them therein.

1.1 Basics of quantum Bose gases

As mentioned in the general introduction, Bose-Einstein condensates (BECs) are useful for producing (comparatively) large-scale quantum objects. In order to do this, they must be made of *bosons*, particles of integer spin that are not prevented from occupying the same quantum state. The atom we are working with is ^{23}Na and it has a nuclear spin of 3/2 and a net electron spin of 1/2. These sum to an integer under any combination of signs, therefore the atom behaves as a boson and does not obey the Pauli exclusion principle, and a number of atoms may share the same state.

Since the phenomenon we are examining involves the particle occupation numbers of different states, we work in the grand canonical ensemble. Let the system have chemical potential μ, and comprise states indexed by k, each state containing N_k particles, each particle having energy E_k. We want to know the occupation number of the k^{th} state and then see under which conditions one of these occupation numbers will grow to be macroscopic. In the grand canonical ensemble, the chemical potential is fixed but the particle number may change. The grand partition function then reads [73, Eqn. 3.2]

$$\mathcal{Z}(\beta,\mu) = \sum_{N_k=0}^{\infty} e^{\beta\mu N_k} Q_{N_k}(\beta), \tag{1.1}$$

where $\beta = 1/k_B T$ with k_B being the Boltzmann constant and T is the temperature of the system, and $Q_{N_k}(\beta) = \sum_k e^{-\beta E_k}$ is the canonical partition function for the k^{th} state.

If one considers the k^{th} state alone, having N_k particles, the grand partition function just for that state is,

$$\mathcal{Z}_k(\beta,\mu) = \sum_{N_k} e^{\beta N_k(\mu - E_k)}. \tag{1.2}$$

For bosons, the sum over N_k must run from zero to infinity (as opposed to $N_k = 0, 1$ for fermions) since the number of particles per state is unbounded. This makes $\mathcal{Z}_k$ the sum over a geometric series with a common ratio of $e^{\beta(\mu - E_k)}$. In order for this to converge, each exponential term must be less than unity, even at $E_k = 0$. This means we must have $\mu < 0$[1] in order to have a system with any well-defined thermodynamic properties at all. The most relevant quantity for us will be the average occupation number when the system is in equilibrium. This can be obtained from the grand partition function [73, Eqn. 3.14],

$$\langle N_k \rangle = \frac{1}{\beta}\frac{\partial}{\partial\mu} \ln \mathcal{Z}_k = \frac{1}{e^{\beta(E_k - \mu)} - 1}, \tag{1.3}$$

where $\langle\cdot\rangle$ indicates the ensemble average. Summing equation (1.3) over k must give the average of the total number of particles in all the states, $\langle N \rangle \equiv \sum_k \langle N_k \rangle$.

When the states are close enough in energy that they can be integrated over, one can make a semiclassical approximation where summing over $\langle N_k \rangle$ should give the same answer as taking the number of particles in each energy state, $\langle N(E) \rangle$; multiplied by the density of states, $\rho(E)$, giving the number of atoms having energy between E and $E + \mathrm{d}E$; and then integrating over the energy. That is to say

$$\langle N \rangle = \sum_{k=0}^{\infty} \frac{1}{e^{\beta(E_k - \mu)} - 1} = \int \rho(E) \langle N(E) \rangle \, \mathrm{d}E \,, \tag{1.4}$$

1. To be precise, $\mu < E_0$, the energy of the ground state, but $E_0 = 0$ is a relevant, limiting case of a free particle [77, §1.1].

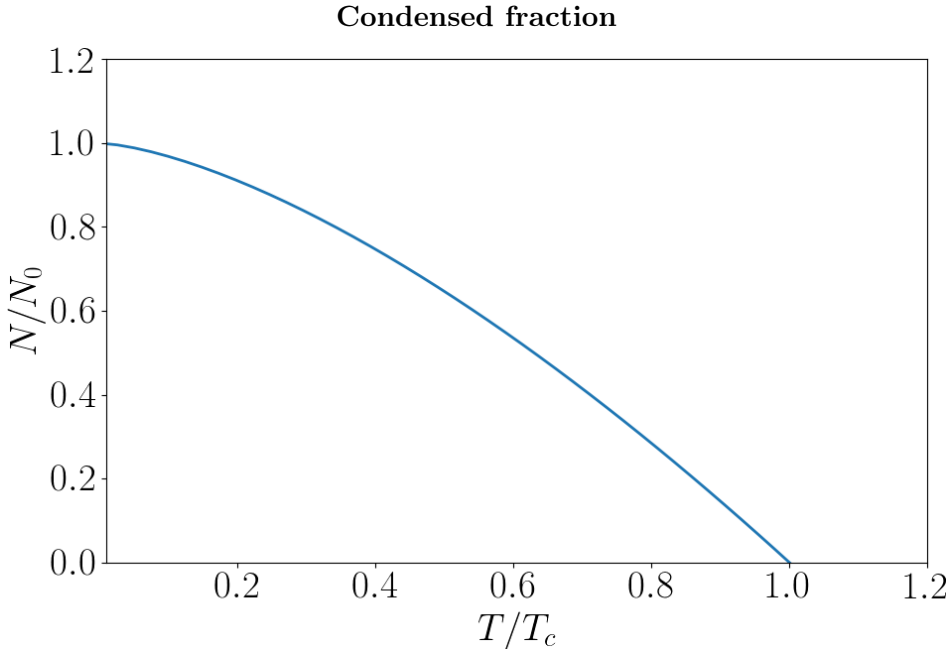

Figure 1.1 — Condensed fraction, N_0/N, as a function of the ratio of the temperature to the critical temperature, T/Tc.

when the states are close in energy. The density of states, $\rho(E)$, will generally depend on the dimension. This will be critical later in Section 1.2 when we want to reduce the dimension below three.

1.1.1 Bose-Einstein condensation in a non-interacting gas

Since bosons can occupy the same quantum state, if there is not enough energy to keep all the particles in excited states — i.e., the excited states become *saturated* containing all the particles they can at that energy — the occupation number of the ground state will start to grow. When the occupation of the ground state becomes macroscopic, this forms a *condensate*. Since there is no interaction energy, only the thermal energy can keep the atoms excited. Therefore, let us write the total average particle number as the sum of the average occupancy of the condensate, $\langle N_0 \rangle$, and of the rest of gas, called the *thermal* component, $\langle N_T \rangle$. That is

$$\langle N \rangle = \langle N_0 \rangle + \langle N_T \rangle = \frac{1}{e^{\beta(E_0-\mu)}-1} + \sum_{k=1}^{\infty} \frac{1}{e^{\beta(E_k-\mu)}-1} \tag{1.5}$$

where the rightmost expression is just an explicit writing-out of the two terms of the middle expression.

For a given, finite temperature, $\langle N_0 \rangle$ will be one or less until $\mu \to E_0^-$, in which case it diverges. This is just the point at which $\langle N_T \rangle$ reaches a finite maximum, $N_{\mathrm{ub}}(T) = \langle N_T \rangle (T, \mu = E_0)$, since $\mu < E_0 < E_{k \geqslant 1}$, ensuring the sum converges. This gives an upper bound on the number of particles that can be in the excited states. If the given temperature is such that $N_{\mathrm{ub}}(T) > \langle N \rangle$, then $\langle N_0 \rangle$ can remain low and the average overall particle number, $\langle N \rangle$, can be distributed among the higher states. When $N_{\mathrm{ub}}(T)$ drops below $\langle N \rangle$, however,

the excited states become saturated and the rest of the particles are forced into the ground state. As the temperature decreases (β rises), so does $N_{\rm ub}$ and particles are driven out of the thermal component and into the condensed component. Thus, $\langle N_0 \rangle$ diverges. The temperature at which particles must start to be forced into the ground state is the *critical temperature*, T_c, defined as the temperature such that [73, Eqn. 3.20]

$$N_{\rm ub}(T_c) = \langle N \rangle \,. \tag{1.6}$$

If the temperature drops below this, particles must start populating the ground state and condensation begins. As the temperature decreases further, $N_{\rm ub}(T)$ decreases further and more particles are forced into the condensate as shown in Figure 1.1.

To see how the density of states will affect $N_{\rm ub}(T)$, let us rewrite $\langle N_k \rangle$ for $k \geq 1$ using the formula for a geometric sum. This reads

$$\langle N_k \rangle = \frac{1}{e^{\beta(E_k - \mu)} - 1} = \sum_{u=1}^{\infty} e^{-u\beta(E_k - \mu)} \quad \text{for} \quad k \geq 1, \tag{1.7}$$

which could be cast as an integral of $\rho(E)$ using a semi-classical approximation.

To say more we need to introduce the specifics of the system so we can talk more precisely about the energy levels and the density of states.

1.1.1.1 Confined particles

Consider particles of mass m moving freely in a box with periodic boundary conditions. Let the side length be L. This makes the momenta $\boldsymbol{p}_k = 2\pi\hbar\boldsymbol{k}/L$, where $\boldsymbol{k}$ is a vector whose components are integers, k_x, k_y, and k_z. The eigenvalues for the energy are $E_k = p_k^2/2m$. Since $\boldsymbol{k}$ can be all zeros, $E_0 = 0$ and μ must be negative. The dispersion relation gives a density of states that is a power law, proportional to $E^{d/2-1}$ where d is the dimension. If the states are close together in energy we may then write $\langle N_T \rangle$ as an integral like

$$\langle N_T \rangle \propto \sum_{k=1}^{\infty} e^{k\beta\mu} \int_0^{\infty} E^{d/2-1} e^{-k\beta E} \mathrm{d}E. \tag{1.8}$$

The integral being summed over is equal to the Γ factorial function [78, Eqn. 6.1.1, 2$^{\rm nd}$ line]

$$\Gamma(z) \equiv s^z \int_0^{\infty} t^{z-1} e^{-st}\, \mathrm{d}t \tag{1.9}$$

with argument $z = d/2$ and an extra factor of $(k\beta)^{-d/2}$ in front to take care of the prefactor of s^z in (1.9). Since $\Gamma(d/2)$ is a constant over the sum we can leave it out and rewrite (1.8) as

$$\langle N_T \rangle \propto \sum_{k=1}^{\infty} \frac{e^{k\beta\mu}}{(\beta k)^{d/2}} = \frac{g_{d/2}(e^{\beta\mu})}{\beta^{d/2}} \tag{1.10}$$

where $g_{d/2}(z)$ is the polylogarithm function of order $d/2$. The polylogarithm of order b is defined for a complex argument, z, as [79, Eqn. 7.1],

$$g_b(z) \equiv \sum_{n=1}^{\infty} \frac{z^n}{n^b}. \tag{1.11}$$

In equation (1.10), $z = e^{\beta\mu}$ is called the *fugacity* [73, p. 17]. Working out the density of states gives the constant of proportionality, we may say [73, Eqn. 3.23]

$$\langle N_T \rangle = \frac{L^3}{\lambda_T^3} g_{d/2}(e^{\beta\mu}) \tag{1.12}$$

where

$$\lambda_T = \sqrt{\frac{2\pi\hbar^2}{mk_BT}} \tag{1.13}$$

is the *thermal de Broglie wavelength* [73, Eqn. 3.24].

To find the upper bound on the number of particles in the excited states, N_{ub}, set μ, $E_0 = 0$. Then the fugacity goes to one, while g_b becomes the Riemann zeta function, $\zeta(b)$ [78, Eqn. 23.2.1]. This means that whether or not condensation occurs depends on the dimension. In three dimensions $g_{3/2}(1) = \zeta(3/2) = 2.612\ldots$, a finite constant. That means by lowering the temperature, λ_T can be increased and the limit on the number of particles in the excited states can be driven down, pushing the rest into the condensate for a given $\langle N \rangle$.

But in two dimensions, the argument of the zeta function, $d/2 = 1$ and so the function diverges [78, Eqn. 23.2.12]. When its argument falls below one, $\zeta(-2 < z < 1) < 0$, meaning that condensation of non-interacting particles in a box of less than three dimensions is impossible at any finite temperature.

However, with different dispersion relations, the density of states can change, changing the exponent of E in equation (1.8) and thereby the order of the polylogarithm, potentially allowing condensation at lower dimension.

Since the trap we will be working with is a harmonic trap, consider the Hamiltonian

$$\hat{H} = -\frac{\hbar}{2m}\nabla^2 + \frac{m}{2}\sum_{i=x,y,z} \omega_i^2 r_i^2 \,. \tag{1.14}$$

For the moment, let $\omega_x = \omega_y = \omega_z = \omega$. The energies are $E_k = \hbar\omega(k_x + k_y + k_z + 3/2)$ and since the energy levels along each direction are evenly spaced by $\hbar\omega$, the density of states is $\rho(E) \propto E^{d-1}/(\hbar\omega)^d \propto E^{d-1}$. Now the polylogarithm function in (1.10) is of order d. This means condensation is possible for dimension $d \geq 2$ in a harmonic trap, the function still being the divergent $g_1(1) = \zeta(1) = \infty$ for one dimension.

The particles in the single-particle ground state are in another phase of matter characterized by an *order parameter*. An order parameter is zero outside of a certain order or phase (as in "phase diagram" or "phase transition" instead of "phase space") of the system and nonzero inside it. They are associated with the presence or absence of some symmetry. When they are non-zero, they are keeping track of some property of the system that was accounted for by a symmetry outside of the phase for which they serve as the order parameter [72, p. 449]. In a condensate, the particles share the same single-particle state, $\Psi_0(\boldsymbol{r})$, and the modulus of this can be taken as the order parameter. The symmetry that is broken inside the condensate is the gauge symmetry associated with the choice in the phase of $\Psi_0(\boldsymbol{r})$. The Hamiltonian (1.14) and all other observables of the system are invariant when multiplied by a constant phase factor. But when a significant number of particles are in the state $\Psi_0(\boldsymbol{r})$ and a value, including a phase, must be chosen for it, this symmetry is broken [73, p. 12]. As we will see later, there can be regimes in a phase diagram where the system will have different characteristics, but an order parameter cannot be defined because no symmetry is broken. We can use the correlations however, as I will discuss in § 1.2.3.

1.1.2 Contact interactions and the Gross–Pitaevskii equation

Now that we understand how condensates can form in gases without interactions we can move closer to our experimental conditions by looking at what happens when there are interactions. If the particles are not in constant contact so that one can consider the particles as undergoing a series of two-body collisions then those collisions can be treated with general scattering theory [71, §122]. This separation is what it means for the gas to be *dilute*, which I will define more precisely in a moment. In quantum scattering theory, all of the results of the interaction are encoded in the far-field behavior of the wave function of a single particle interacting with an effective potential which encodes their relative motion. The goal then is to calculate the scattering amplitude for all angles giving the probability that the particle will scatter into that angle. At low relative momenta,[2] $p_r \to 0$, the scattering amplitude converges to a single value for all angles [73, §9.2], equal in magnitude to a quantity called the *s-wave scattering length* or just *scattering length*, denoted a. In this case, the effective potential in the scattering problem is a hard sphere of radius a [71, p. 502]. Since the scattering amplitude as a function of the scattering angle encodes everything about the collision, and in the condensate all of the particles have momenta that are very low, the scattering amplitude converges to a value independent of the scattering angle. Then everything about the interactions in the condensate should be able to be described in terms of the scattering length.

Although particles can interact in any number of ways, (including some that lead to exactly solvable systems in one-dimension, see [80] and references therein.), our neutral sodium atoms do not affect each other at long ranges. Thus we only need consider *contact interactions*, that is that the "far field" in the scattering problem is any non-vanishing distance. That means the interactions can be described by a delta function

$$g_{3\mathrm{D}}\delta(\boldsymbol{r}-\boldsymbol{r}') \tag{1.15}$$

where the $g_{3\mathrm{D}}$ parameter[3] is the amplitude of the delta function and so is called the three-dimensional *interaction strength*.

Now I can be more precise about what it means to be "dilute." It means that the scattering length is much less than the average inter-particle distance, that is

$$|a| \ll n^{-1/d}. \tag{1.16}$$

The absolute value is necessary since scattering length can be positive or negative, depending on the phase shift imparted by the interaction. Negative scattering lengths are typically associate with attractive interactions. Since all properties of the interaction should be expressible in terms of a, we can write $g_{3\mathrm{D}}$ as [73, Eqn. 4.11]

$$g_{3\mathrm{D}} = \frac{4\pi\hbar^2 a}{m}. \tag{1.17}$$

Now that we have the interaction strength, it would be nice to have a way to describe the interacting gas in a way that does not involve just summing over all the particle pairs,

2. This is when λ_T is large compared to the distance of the inter-atomic interaction. Taking the Bohr radius, $a_0 \sim 10^{-10}\,\mathrm{m}$ to give this length scale, this happens when $T \lesssim 1\,\mathrm{K}$.
3. I use three different gs to be consistent with the literature. Those with a numerical subscript $g_n(z)$, are polylogarithmic functions of order n; those with a subscript like $g_{n\mathrm{D}}$ are real numbers giving the n-dimensional interaction strength; and those with a parenthetical superscript, $g^{(n)}$, are correlation functions of order n.

as this would quickly become intractable — it would be worse than an Ising model where only neighbors can interact. Ising models have so far resisted analytic solutions in three dimensions [81]. Instead the gas is treated with a *mean-field* method, one of replacing the quantum operator, $\hat{\Psi}(\boldsymbol{r},t)$, with a classical field, the wave function of the condensate, $\Psi_0(\boldsymbol{r},t)$, and neglecting particle correlations. This is why mean-field theories break down in strongly-interacting systems: strong interactions lead to strong correlations between the particles. For example, if the particles have very strong repulsive interactions, finding a particle at a position excludes the possibility of finding another particle at that same position. So finding a particle gives definite information about the location of other particles and so they are strongly correlated and neglecting correlations will lead to poor approximations. Neglecting these correlations leads to the Hartree–Fock approximation of the many-body wave function, where it is written as a product of a single single-particle wave function at different locations. This reads

$$\Phi_0(\boldsymbol{r}_1,\boldsymbol{r}_2,\ldots,\boldsymbol{r}_N)=\left(\frac{\Psi_0(\boldsymbol{r}_1)}{\sqrt{N}}\right)\left(\frac{\Psi_0(\boldsymbol{r}_2)}{\sqrt{N}}\right)\ldots\left(\frac{\Psi_0(\boldsymbol{r}_N)}{\sqrt{N}}\right). \tag{1.18}$$

This should work to describe the condensate. As long as the condensed fraction is significant and interactions are weak, and we may say $\Psi_0=\langle\hat{\Psi}\rangle$, this should in turn provide a good description of the overall gas. As long as the external, confining potential does not vary significantly over the range of inter-particle forces, it will not appreciably affect the interactions or interaction energy, and the assumptions that lead to $g_{3\mathrm{D}}$ will still hold. Then, for the condensate at any point, the interaction energy should just depend on the interaction strength and the field density, $|\Psi_0(\boldsymbol{r},t)|^2$. This density times the interaction strength will give the interaction energy. In an external potential $V(\boldsymbol{r},t)$, the Hamiltonian will then be [82,83]

$$\hat{H}=\frac{-\hbar^2}{2m}\nabla^2+V(\boldsymbol{r},t)+g_{3\mathrm{D}}\left|\Psi_0(\boldsymbol{r},t)\right|^2. \tag{1.19}$$

Considering stationary states, $\Psi_0(\boldsymbol{r},t)=\Psi_0(\boldsymbol{r})e^{-i\mu t/\hbar}$ [73, Eqn. 5.16], (1.19) yields the time-independent Gross–Pitaevskii equation [73, Eqn. 5.18]

$$\left[\frac{-\hbar^2}{2m}\nabla^2+V(\boldsymbol{r})+g_{3\mathrm{D}}\left|\Psi_0(\boldsymbol{r})\right|^2\right]\Psi_0(\boldsymbol{r})=\mu\Psi_0(\boldsymbol{r}), \tag{1.20}$$

where $|\Psi_0(\boldsymbol{r})|^2=n(\boldsymbol{r})$, the density of the condensate at $\boldsymbol{r}$. The appearance of the chemical potential μ is not surprising if one remembers that the condensate is an eigenstate of the single-particle Hamiltonian and if the energy of a particle in the condensate changes, it must change the energy of the condensate by μ. The particle number and the chemical potential are then related by $\mu=\partial E/\partial N$ where N is the total particle number and E is given by the energy functional [73, Eqn. 5.8]

$$E=\int\left[\frac{\hbar^2}{2m}|\boldsymbol{\nabla}\Psi_0|^2+V(\boldsymbol{r})|\Psi_0|^2+\frac{g_{3\mathrm{D}}}{2}|\Psi_0|^4\right]\mathrm{d}\boldsymbol{r}. \tag{1.21}$$

Then fixing either μ or N, will fix the other and we have an equation we can solve for the wave function of the condensate.

Having a term in the Hamiltonian that depends on the density of the condensate makes it natural to ask over what length scale the density can be expected to change. This is called the *healing length*, ξ, which reads [73, Eqn.4.37]

$$\xi=\frac{\hbar}{\sqrt{2mng_{3\mathrm{D}}}}. \tag{1.22}$$

This is the length scale associated with the point at which the excitations of the gas reach the energy scale of the interaction energy. Bogoliubov theory describes the low-energy excitations of a weakly-interacting condensate as non-interacting quasi-particles [84]. The dispersion relation of these excitations is [73, 4.31]

$$\epsilon(p) = \sqrt{\frac{g_{3\mathrm{D}}n}{m}p^2 + \left(\frac{p^2}{2m}\right)^2}. \tag{1.23}$$

There is a linear part for low-momentum excitations, and a quadratic part for higher-energy ones which have a particle-like $\epsilon(p) = p^2/2m$ dispersion relation. The latter should dominate over the former when $p^2/2m \sim g_{3\mathrm{D}}n$. Setting $p^2/2m = g_{3\mathrm{D}}n$ and associating a length scale with this momentum, $p = \hbar/\xi$, one obtains (1.22).

It should be noted that equation (1.20) is based on a number of assumptions and approximations and only describes the behavior of the condensate, neglecting the particles in excited states. If a significant part of the gas is still excited, i.e. if N/N_0 differs significantly from one, then interactions with this thermal part will be significant and neglected by (1.20). It also relies on there being many particles in the ground state to average over so only applies if N is large. It also will break down if the interactions are too strong since this will cause correlations. What is more, if the scattering length is on the order of the inter-particle distance, the diltue condition (1.16) breaks down and equation (1.20) cannot be used. I will discuss further differences in regimes in § 1.2.2 and following.

Now we have what we need to start talking in more detail about one-dimensional systems with interactions.

1.2 One-dimensional systems

The basic trick behind making a one-dimensional system in the real, three-dimensional world is to make excitations in two of the dimensions so energetically costly, and make the overall energy so low, that there is not enough energy in the entire system to excite even one transverse mode. This way, all excitations are in the *longitudinal* mode, *transverse* dynamics are prevented, particles cannot get past each other without interacting, and the system is effectively one dimensional. This was first realized experimentally by Görlitz et al. [85], although theoretical research into one-dimensional Bose gases goes back much further [48, 86, 87]. Even before this "freezing out" of the dynamics in two dimensions there are regimes where phase fluctuations start to dominate in elongated traps [88] giving a picture of what will happen in the one-dimensional regime even before it is reached.

Because, as mentioned above, there is no condensation in one dimension at the thermodynamic limit, so there can be no true phase transitions in one-dimensional systems. This means we cannot break the phase diagram up into different phases via an order parameter. However there can be different regions of the phase diagram where different features dominate (e.g. kinetic vs. interaction energy) and the system has different behaviors. These can be characterized by differences in the fluctuations in different regions.

While one dimensional systems are amenable to approximate techniques like Bogoliubov and mean-field theory where they are applicable, they have the attractive feature of admitting exact solutions. Apart from the theoretical interest, they offer experimentalists excellent points of comparison. Precise predictions for experimentally testable quantities can be made and properties of the system can be extracted by fits to analytical formulae.

Transverse energy scales in the effectively 1-D harmonic trap

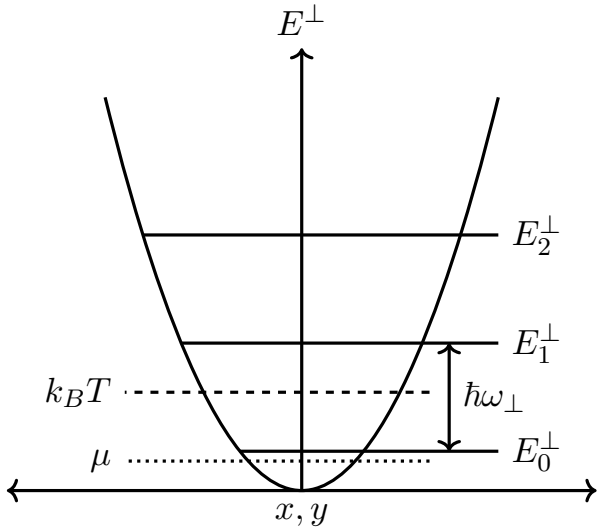

Figure 1.2 — Energy given to the atoms by the transverse confinement, $E^\perp$, where $E_n^\perp$ is the n^{th} energy level in the transverse harmonic trap. This is the non-interacting case so $\mu < E_0^\perp$. With even weak interactions, μ would rise above this level (1.28).

First I will explain in more detail how to make an effectively one-dimensional system and relate its interaction strength to the more general and well-studied scattering length, a. Then I will need to establish a way to tell different regimes of the gas apart. Then I will discuss some of those means uniquely available in one dimension: exact solutions for the wave function and properties of the gas. In each of these sections I will point out crossovers they predict. Then I will sketch an outline of the phase diagram we want to explore in our experiment, giving the characteristics of its different regions.

1.2.1 Reducing the dimension

The method for producing a gas that behaves as if it is one-dimensional is to confine it with an external potential so strongly in the other two dimensions, and reduce the energy in the system so much, that excitations in the transverse direction are inaccessible [89]. In the experiments that have made one-dimensional gases, from the first experiment [85] to ones at weak [90] or variable [91,92] effective interaction strengths all use harmonic transverse trapping potentials. This has another advantage of making the ground state in the transverse direction a Gaussian and an effective, one-dimensional scattering length can be calculated from the three dimensional scattering length [93]. So I will discuss this case where the longitudinal potential is homogeneous, but the transverse confining potential is harmonic.

The energetic scheme for confinement is illustrated in Figure 1.2. The transverse potential is harmonic in both directions, x and y, with frequency $\omega_\perp$, making the transverse energy levels $E_n^\perp = (n + 1/2)\hbar\omega_\perp$ where $n = 0, 1, \ldots$. If $k_BT < E_1^\perp$, there will not be enough kinetic energy in the system to form any transverse excitations. And if $\mu < E_1^\perp$ then a particle entering or leaving the condensate will also not impart enough energy to excite the first excited transverse state. In this way we can be sure that all the atoms are in the ground state in the transverse direction and that all the dynamics are happening in the longitudinal

direction z. In practice, temperatures this low are not always achieved with the temperature ranging from $k_BT \simeq .9\hbar\omega_\perp$ [91] or $1.45\hbar\omega_\perp$ [92], which would keep the thermal energy below the first excited state with $E_1^\perp = 1.5\hbar\omega_\perp$, to $kk_BT = 2.4\hbar\omega_\perp$ [91] or $2.75\hbar\omega_\perp$ [94]. This means that some fraction of the atoms still have three-dimensional dynamics, but certain features of one-dimensional systems can be present as in [92] if the longitudinal dynamics are dominant.

Since higher transverse frequencies mean the one-dimensional regime can be reached at higher temperature, $\omega_\perp/2\pi \lessapprox 3\,\mathrm{kHz}$ are typical [91,92,94] but ones as low as $\omega_\perp/2\pi \approx 360\,\mathrm{Hz}$ have been used [85].

It would be conceptually useful to have a one-dimensional Hamiltonian; one where all of the motion is in one dimension, z, and the interactions are still contact interactions. This is called the Lieb–Liniger Hamiltonian [48]. For N particles, it has the form

$$\hat{H} = -\frac{\hbar^2}{2m}\sum_{i=1}^{N}\frac{\partial^2}{\partial z^2} + g_{1\mathrm{D}}\sum_{\{i,j\}}\delta(z_i - z_j) \tag{1.24}$$

where δ is the Dirac delta-function, $\{i,j\}$ represents distinct particle pairs, and $g_{1\mathrm{D}}$ is a one-dimensional interaction strength, analogous to $g_{3\mathrm{D}}$ in 1.20, that we must relate to the more well-measured three-dimensional interaction properties, i.e. the scattering length, a. To do this we will make a one-dimensional version of the Gross-Pitaevskii equation (1.20) for the cylinder trap and then extend it to an elongated harmonic trap.

As mentioned above, in a cylindrical trap, where the transverse confinement is harmonic but the longitudinal one is homogeneous, it is in the transverse direction that the system is expected to appear in the ground state. Thus, it is possible to rewrite Ψ_0 as $\Psi_0 = \sqrt{n_{1\mathrm{D}}}f(\rho_\perp)/a_\perp$, a function of the one-dimensional radius, $n_{1\mathrm{D}} = N/L$ where N the total atom number and L is the length of the trap; transverse radius, $\rho_\perp = r_\perp/a_\perp$, where $r_\perp$ is just the radial distance from the central axis; and $a_\perp = \sqrt{\hbar/m\omega_\perp}$ is the length scale associated with transverse harmonic oscillator. It then becomes possible to write a one-dimensional version of 1.20 as in terms of this function, $f(\rho_\perp)$. This reads [73, Eqn. 17.3]

$$\left[-\frac{1}{2}\frac{\partial^2}{\partial\rho_\perp^2} - \frac{1}{2\rho_\perp}\frac{\partial}{\partial\rho_\perp} + \frac{1}{2}\rho_\perp^2 + 4\pi a n_{1\mathrm{D}} f^2\right] f(\rho_\perp) = \frac{\mu}{\hbar\omega_\perp}f(\rho_\perp) \tag{1.25}$$

with the normalization condition that

$$2\pi\int \rho_\perp |f(\rho_\perp)|^2\,\mathrm{d}\rho_\perp = 1. \tag{1.26}$$

From here we can get the limit where we expect the cross-over to the one-dimensional regime and from there the chemical potential in those regimes. Then we can get the effective one-dimensional interaction strength with which particles moving in the longitudinal direction will interact, since this is not the same as $g_{3\mathrm{D}}$.

When $an_{1\mathrm{D}} \gg 1$, the interaction energy in (1.25) will dominate and the system is in the Thomas–Fermi regime in the radial direction.[4] Here, excitations of the transverse harmonic oscillator are expected and the transverse density should take the form of a parabola (the Thomas–Fermi profile). Since interactions dominate, in the ground state, the kinetic energy term in (1.25) can be neglected. Using this and (1.26) one can find the chemical potential

4. The Thomas–Fermi approximation is a more general construction with broader applicability then one-dimensional systems [71, §70], even within cold atom physics [73, § 5.2].

where the system is expected to cross between the one-dimensional regime and the Thomas–Fermi one. This is just is [73, Eqn. 17.4]

$$\frac{\mu}{\hbar\omega_\perp} = 2\sqrt{an_{1\mathrm{D}}}. \tag{1.27}$$

In the case where $an_{1\mathrm{D}} \ll 1$, the transverse profile will be in the Gaussian ground state of the harmonic oscillator, and we are in a weakly interacting one-dimensional system. The chemical potential is then [73, Eqn. 17.6]

$$\frac{\mu}{\hbar\omega_\perp} = 1 + 2an_{1\mathrm{D}}. \tag{1.28}$$

The factor of unity on the right-hand side is from the transverse ground state. If one wants to define a $\mu_{1\mathrm{D}}$ for the one-dimensional system, this can be omitted.

By separation of variables into axial and radial components in the harmonic trap and breaking the wave function up into odd and even components, it is possible to define a one-dimensional effective scattering length in terms of the three-dimensional scattering length [73, Eqn. 17.35],

$$a_{1\mathrm{D}} = \frac{a_\perp^2}{a}. \tag{1.29}$$

This behaves as the three dimensional scattering length, giving the phase shift in the wave scattered by a hard-sphere potential and giving a way to define the interaction parameters. The one-dimensional interaction strength can then be defined in terms of the one-dimensional, or three-dimensional, scattering length as [73, p. 332]

$$g_{1\mathrm{D}} = \frac{2\hbar^2}{ma_{1\mathrm{D}}} = \frac{2\hbar^2}{m}\frac{a}{a_\perp^2}. \tag{1.30}$$

See [93] for details of the calculation of the two quantities above.

Because these one-dimensional interaction parameters are derived from the Gross-Pitaevskii equation, which relies on mean-field theory, the above forms require that $a \ll a_\perp$. If this condition is not met, they must be renormalized [93].

Adding a harmonic confinement in the longitudinal direction, $V(z) = m\omega_z^2 z^2/2$, means the density, and hence the interaction energy, will no longer be uniform in the longitudinal direction. Instead, the one-dimensional density will be function of the longitudinal position, $n_{1\mathrm{D}} = n_{1\mathrm{D}}(z)$ and if the longitudinal trapping potential is still smooth, we can make a *local density approximation* for the chemical potential,

$$\mu(n_{1\mathrm{D}}(z)) + V(z) = \mu_0 \tag{1.31}$$

where μ_0 is the overall chemical potential which is still fixed by insisting that $\int n_{1\mathrm{D}}(z)\,\mathrm{d}z = N$, where N is the total atom number. This means we assume that the density, and properties dependent on it, are fixed at any particular $z + \mathrm{d}z$, but may vary with z. The transition between the radial Thomas-Fermi regime and being in the radial ground state then depends on whether the quantity [73, Eqn. 17.8]

$$N\frac{\omega_z}{\omega_\perp}\frac{a}{a_\perp} \tag{1.32}$$

is large or small compared to unity. If $N\omega_z a/\omega_\perp a_\perp \gg 1$ then the system is in the Thomas–Fermi regime and if $N\omega_z a/\omega_\perp a_\perp \ll 1$ then it is the radial ground state.

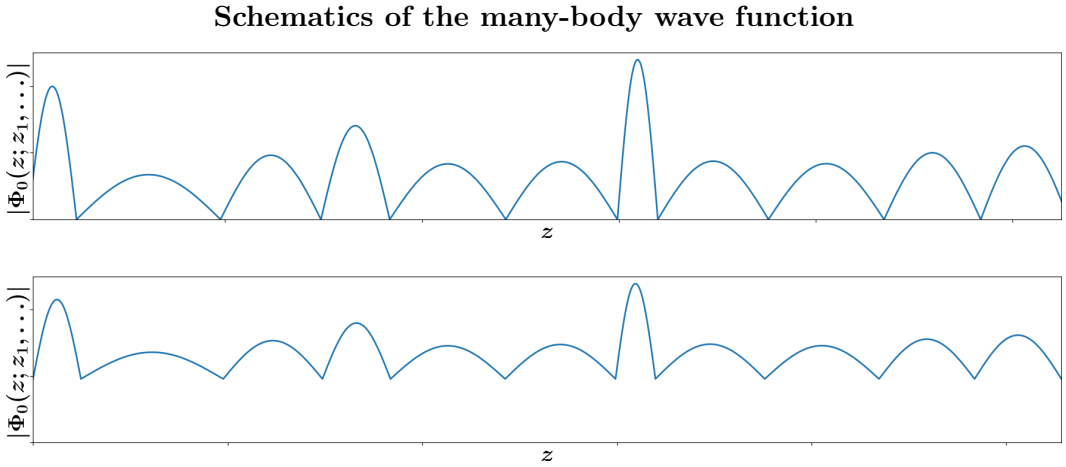

Figure 1.3 — Top — Absolute value of the one-dimensional many-body wave function with all but one particle fixed, $\Phi_0(z; z_1, z_2, \ldots, z_N)$ (1.18), in the strongly-interacting regime. Here, $\Phi_0 \to 0$ at the location of the atoms, the wave function is strongly bent and the healing length (1.22) is shortened. Bottom — Similar in the weakly-interacting regime. Here Φ_0 need not reach zero and the healing length is longer. Adapted from [47, Fig. 2]

These, and (1.25) tell us that in order to keep the chemical potential low enough to be in the one-dimensional regime, we want the energy associated with interactions, an_{1D} to be small compared to the energy associated with the transverse confinement, $\hbar\omega_\perp$. This can occur when the interactions are weak, that is a is small, but this limits the systems one can explore and exclude some exactly solvable systems [69] whose solutions rely on the atoms being impenetrable. So, the usual thing is to work in very elongated traps to keep $\omega_z/\omega_\perp$ low. Typically this means a confining frequency aspect ratios of $\omega_z/\omega_\perp \sim 1/100$ [47, 85, 91]. For comparison, a sodium atom in a harmonic trap of $\omega_\perp = 2\,\mathrm{kHz}$ will yield the ratio $a/a_\perp \approx 1/500$ and a typical atom number for the one-dimensional regime would be $N \sim 10^4$ [85].

1.2.2 One-dimensional interaction regimes

From the above discussions it should be clear that much depends on the energy associated with interactions, which scales like $g_{1D}n_{1D}$, which gives the interaction energy per unit length and so depends both on the interaction strength and the density. Since there is no true condensate in one dimension except at g_{3D} and $T = 0$, there will always be excitations of some kind. Whether these excitations are due to kinetic or interaction energy is the distinguishing feature between the *strongly-* and *weakly-interacting regimes*.

1.2.2.1 Strong vs. weak interactions

In one dimension the average inter-particle separation will be by definition $1/n_{1D}$. Since kinetic energy is given by $-\hbar^2/2m \cdot \partial^2\psi/\partial z^2$, it is greater the more steeply bent the wave function is and will also depend on the density. At the points where the atoms collide, $|\Psi|$ will reach a minimum. If the particles are hard-cored, meaning $g_{1D} \to \infty$, then $|\Psi| \to 0$ where the particles meet, as in the upper panel of Figure 1.3. For finite values of g_{1D}, the minimum will be higher as in the lower panel. For a given interaction strength, the average kinetic energy in a given length will also depend on the density; the "valleys" caused by the particles become both more frequent, and steeper, as the number of particles per length increases. Consider a

the example of the ground state between two hard-core particles separated by a length $1/n_{1D}$. The ground state wave function between the two particles will be that of a particle in a box of length $1/n_{1D}$. That is, $\psi(z) \propto \sin(\pi z n_{1D})$. The kinetic energy operator will then yield

$$\frac{-\hbar^2}{2m}\frac{\partial^2\psi}{\partial z^2} \propto \frac{\hbar^2}{2m} n_{1D}^2 \sin(\pi z n_{1D}). \tag{1.33}$$

Therefore, the kinetic energy scale for the ground state should go like the linear density squared [48, Eq. 3.] and we can associate an energy scale with the kinetic energy of $E_{kin} = \hbar^2 n_{1D}^2/2m$. When one of either E_{kin} or $g_{1D}n_{1D}$ dominates over the other, we can say whether we are in a regime dominated by interactions or not. So let their ratio be the dimensionless parameter

$$\gamma = \frac{m g_{1D}}{\hbar^2 n_{1D}}, \tag{1.34}$$

where n_{1D} will depend on z if the longitudinal trapping is inhomogeneous.

This γ is the Lieb–Liniger parameter from [48] used to characterize the strength of the interactions. When $\gamma \ll 1$ we can be sure that the kinetic energy will dominate over the interaction energy and we are in the weakly-interacting regime (Figure 1.3, below) where the Gross–Pitaevskii equation and mean-field theories can be used. When $\gamma \gg 1$ we can be sure that interaction energy will dominate and we are in the strongly-interacting regime (Figure 1.3 above) where the interaction is strong enough to act like the Pauli exclusion principle does for fermions. Here the particle locations are correlated, they cannot be found in the same place and tend not to be found near each other. Therefore, the Hartree–Fock approximation (1.18) breaks down and the Gross–Pitaevskii equation does not hold but exact solutions become tractable [95].

I will go into more detail about these regimes in § 1.2.4. Because they involve systems with significant interaction energies, mean-field theory will be insufficient to discuss them all. Fortunately in one dimension, we have exact solutions we can appeal to that are valid for strong interactions.

1.2.2.2 Strong interactions and exact solutions

One dimensional systems are unique in the study of quantum many-body problems in that they admit exact solutions (see [80] and references therein). These are theoretically gratifying and serve as points of departure for other theories and approximations to widen their scope. The importance for our experiment is that they make some predictions that do not arise in approximate theories like Bogoliubov theory and the local density approximation.

The exact solutions I will be talking about will be ones for systems interacting via contact interactions although other exact models do exist [96]. First I will treat systems at zero temperature with the Lieb and Liniger model [48] which makes use of the Bethe ansatz. Then I will touch on the Yang-Yang generalization of this [50] with gives the thermodynamic properties of a system at finite temperature.

At zero temperature: the Lieb and Liniger model Three years after Girardau solved the problem for impenetrable bosons by mapping them to fermions in 1960 [86], Lieb and Liniger [48] found an exact solution for systems of particles with repulsive interactions, in a periodic box, obeying the Hamiltonian (1.24).

Their method relies on the *Bethe ansatz* [68, 97]. The essence of the ansatz is to assume that the wave function, Ψ, is a sum of plane waves. It must be the case that Ψ is symmetric

or anti-symmetric under exchange of particle labels if the particles are bosons or fermions, respectively. This lets one turn a single valid sum of plane waves, corresponding to a certain ordering of particles, into the whole set of plane waves that must be summed over by permuting over all of the possible particle orderings. In a closed system, such as they treated, this means the action of the entire permutation group and includes every possible ordering. In an open system, such as those solved by J. B. McGuire [87] and later by B. Sutherland [98] and M. Olshanii [99], the system undergoes a particular set of collisions defining a particular set of permutations.[5]

The solutions obtained via the Bethe ansatz method have the form

$$\Psi(x_1, x_2, \ldots, x_N) = \sum_{\{\hat{\mathcal{P}}\}} b(\hat{\mathcal{P}})\hat{\mathcal{P}} \exp\left(i \sum_{j=1}^{N} k_j x_j \right) \tag{1.35}$$

where $\{\hat{\mathcal{P}}\}$ is the set of operators that each give a different permutation of the particles; $\{k\}$ is the set of quasi-momenta that respect the relationship

$$\frac{\hbar^2}{2m} \sum_j k_j^2 = E \tag{1.36}$$

and $\{b\}$ is a set of coefficients, one fore each permutation, that must be solved for. This can be done by turning the δ-potential interaction into boundary conditions on Ψ; namely that Ψ is continuous but there is a jump of size $g_{1\mathrm{D}}$ in the derivative whenever particles collide. This, the symmetry conditions on Ψ under exchange of particles, and (1.36) yield a set of parallel equations large enough to solve for all b values and construct the exact solution (1.35).

The spectrum predicted by Lieb in a second, concurrent paper [49], is shown in Figure 1.4. It predicts a *second branch* in the excitation spectrum, in addition to the excitations predicted by Bogoliubov theory. This *Lieb II branch* comprises the "hole excitations", so called because they are gaps in the energy spectrum. The Bogoliubov excitations are also predicted as the *Lieb I branch* "particle excitations", not to be confused with the higher-momentum, particle-like part of the Bogoliubov spectrum. These spectra are compared in Figure 1.4.

Both the Lieb I excitations [101] and the Lieb II excitations [102] have different spectra in different regimes of the zero temperature gas, defined by momentum of the excitations and the interaction strength. These spectra can be probed via Bragg spectroscopy which I will outline in Chapter 2.

This solution holds at any non-negative interaction strength, it is only valid at zero temperature, when all of the kinetic energy is contained in coherent excitations.

At finite temperature: the Yang-Yang model In 1969 Yang and Yang extended the applicability of this model to systems of finite temperature at thermal equilibrium [50]. They proved that the equations given by Bethe's ansatz are applicable at all temperatures and so wave functions constructed by that method are valid at finite temperature. This method also gives the thermodynamic properties of the gas at equilibrium and the excitation spectra at finite temperature.

This model gives different predictions for the behavior of a gas at finite temperature than the local density approximation, which have been verified by experiment [91]. Starting from the Yang-Yang results for a gas at equilibrium, the extension to the dynamics of out-of

5. I will come back to this briefly in 6.2.2.2 since billiards provide a method of generating these permutations.

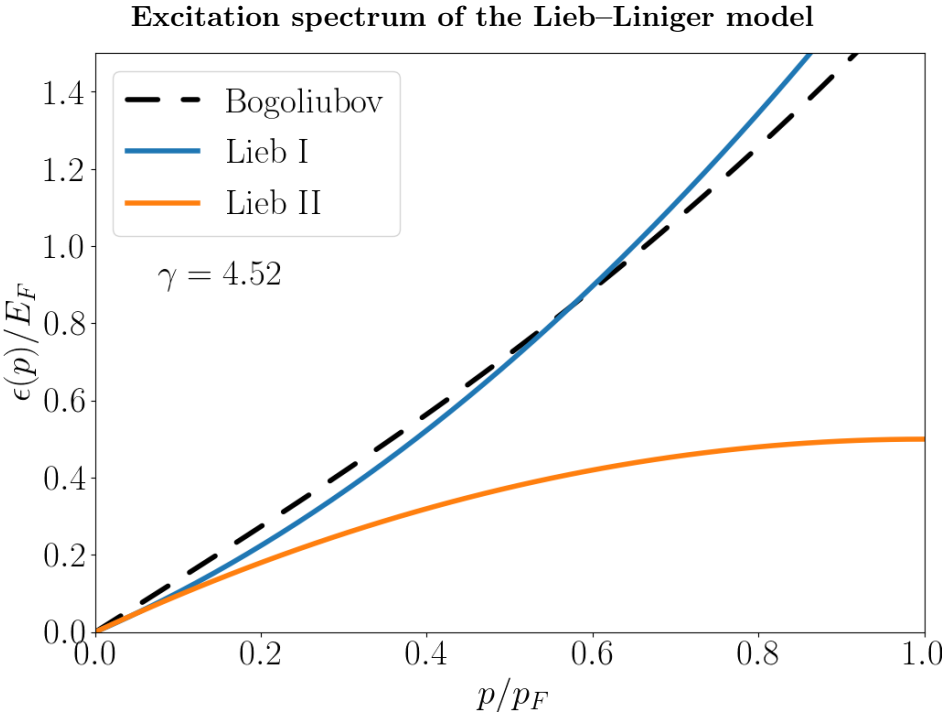

Figure 1.4 — Excitation spectrum predicted by the Lieb and Liniger solution (solid lines) compared to the Bogoliubov mean-field spectrum (equation (1.23), dashed line). Notice that the Lieb and Liniger solution predicts a second branch of excitations. These are associated with "hole excitations." The momentum and energy have been scaled by the Fermi momentum, $p_F = \pi\hbar n_{1D}$, and the Fermi energy, $E_F = p_F^2/2m$ respectively [73, p. 334]. Adapted from [100, Fig. 24.6].

equilibrium systems was made, going under the name *generalized hydrodynamics* [103, 104] which also seems to give closer description of experimental results than the local density approximation and mean-field theories [105].

1.2.3 Correlations and fluctuations

Beyond the division between strong and weak interactions, the phase diagram of a one-dimensional gas can be broken up into smaller regions. These will exhibit differences in their *correlation function* or *degree of coherence*. This notion arose in the context of quantum optics [106] and these functions are proportional to the average correlations within (or between) a field (or fields). They are also perfectly well defined for quantum field operators such as $\hat{\Psi}_k = \hat{a}_k\psi_k$, where $\hat{a}_k$ is the annihilation operator for a particle in state k and ψ_k is that state's wave function. In general, they can measure correlations in space and time. I will only consider stationary states so I drop the time dependence. As we will see, these provide a method for distinguishing different regions of the phase diagram that is more granular than an order parameter. An order parameter will distinguish between truly different phases of a system, characterized by a broken symmetry but there can exist different regimes with different physical properties that will have different correlation characteristics even if they have the same symmetries.

I will discuss the first two degrees of coherence, the first being associated with correla-

tions in the phase fluctuations and the second being associated with correlations in density fluctuations.

1.2.3.1 First-order correlation function

Let $\hat{\Psi}_k(\boldsymbol{r})$ be the field operator for state k. This lets us define the average number density $n(\boldsymbol{r}) = \langle\hat{\Psi}^\dagger(\boldsymbol{r})\hat{\Psi}(\boldsymbol{r})\rangle$ where $\hat{\Psi}(\boldsymbol{r}) = \sum_k \hat{a}_k\psi_k(\boldsymbol{r})/\sqrt{V}$ is the field operator in the second-quantization picture for a system of volume V. Then the *first-order correlation function* may be written [107, Eqn. 2.16]

$$g^{(1)}(\boldsymbol{r},\boldsymbol{r}') = \frac{\langle\hat{\Psi}^\dagger(\boldsymbol{r})\hat{\Psi}(\boldsymbol{r}')\rangle}{\sqrt{n(\boldsymbol{r})n(\boldsymbol{r}')}} \tag{1.37}$$

where the average is taken over the position, $\boldsymbol{r},\boldsymbol{r}'$. This measures phase correlations since the average density will be normalized off, leaving only the overlap of the phases.

It is common to also define $g^{(1)}(\boldsymbol{r})$ where $\boldsymbol{r}' = \boldsymbol{r}_0$, the position of maximum density, serves as a fixed point of comparison and is left out of the argument.

Far from degeneracy, in the cylinder trap, with no interactions, $E_0 - \mu \gg k_BT$, making $\langle N_T\rangle \ll N_{\text{ub}}$, or equivalently, the de Broglie wavelength is much smaller than even the smallest average inter-particle distance, $\lambda_T \ll 1/n$. Here the correlations will be on the scale of the de Broglie wavelength and (1.37) can be approximated as [47, Eqn. 17]

$$g^{(1)}(z) \simeq e^{-\delta z^2/4\pi\lambda_T^2}, \tag{1.38}$$

where $\delta z = |z - z_0|$ and z is the long axis of the cylindrical potential. In the degenerate regime, $E_0 - \mu \ll k_BT$, the fugacity in (1.12) goes to one, and many particles are found within a de Broglie wave length from each other, i.e. $\lambda_T \gg 1/n$. This puts the correlation length on the scale of $n\lambda_T^2$. Now (1.37) is approximately [47, Eqn. 18]

$$g^{(1)}(z) \simeq e^{-2\pi\delta z/n\lambda_T^2}. \tag{1.39}$$

In one dimension, even at $T = 0$, N_{ub} is divergent and so particles can always populate the excited states, if they can find the energy. At zero temperature, interaction energy can be enough to kick particles from the ground state into the excited states. This means for systems of finite temperature or with appreciable interactions there will always be populations in multiple states of different phases, which will be proportional to the ground-state population even in the thermodynamic limit. Then there will be significant fluctuations in the phase [108] and phase correlations will fall off exponentially [109], as in (1.39), but the correlation length will no longer be the de Broglie wavelength but ℓ_ϕ, some function of the temperature and the chemical potential [88]. Then (1.39) has a more general form and reads

$$g^{(1)}(z) \sim e^{\frac{-\delta z}{\ell_\phi}}. \tag{1.40}$$

1.2.3.2 Second-order correlation function

Once can define a *second-order* or *two-point* correlation function [107, Eqn. 2.17],

$$g^{(2)}(\boldsymbol{r},\boldsymbol{r}') = \frac{\langle\hat{\Psi}^\dagger(\boldsymbol{r}')\hat{\Psi}^\dagger(\boldsymbol{r})\hat{\Psi}(\boldsymbol{r})\hat{\Psi}(\boldsymbol{r}')\rangle}{n(\boldsymbol{r})n(\boldsymbol{r}')}. \tag{1.41}$$

This is proportional to finding an actual particle at $\boldsymbol{r}'$ if one has already found one at $\boldsymbol{r}$; it measures the correlations in the fluctuations around the density at $\boldsymbol{r}$.

Let me rewrite (1.41) in the form used by Kheruntsyan et al. in [56] from whom I take my definition of the phase regimes. This will also make the connection to density fluctuations clearer. This will be for a one-dimensional gas with the long dimension along z. The field operator $\hat{\Psi}(z)$ has the commutation relation [71, Eqn. 64.22]

$$\left[\hat{\Psi}(z), \hat{\Psi}^\dagger(z')\right] = \delta(z - z'). \tag{1.42}$$

Therefore at $z' = z$, (1.41) reads

$$g^{(2)}(z, z) = \frac{\left\langle \hat{\Psi}^\dagger \hat{\Psi} \hat{\Psi}^\dagger \hat{\Psi} \right\rangle - \left\langle \hat{\Psi}^\dagger \hat{\Psi} \right\rangle}{n^2}, \tag{1.43}$$

where the argument of $\hat{\Psi}$ has been dropped for neatness. Defining the variance of the density as $\left\langle \delta n^2(z) \right\rangle = \left\langle n^2(z) \right\rangle - \langle n(z) \rangle^2$ and inserting $0 = \langle n(z) \rangle^2 - \langle n(z) \rangle^2$ into equation (1.43) yields

$$g^{(2)}(z) \equiv g^{(2)}(z, z) = 1 + \frac{\left\langle \delta n^2(z) \right\rangle - \langle n(z) \rangle}{n(z)^2}. \tag{1.44}$$

It is this form I will use to distinguish the regions of the phase diagram, in line with [56]. When z is set to zero, where the density is maximum, this becomes $g^{(2)}(0)$ and it is this we will use to distinguish phase regions.

This correlation function can be calculated empirically by measuring fluctuations and calculating (1.44). This is a key consideration in our experimental design. In order to say things about correlations and fluctuations, statistics must be gathered over a significant number of experimental realizations. The higher the repetition rate on the experiment is, the more instances can be averaged over and the more reliable the statistics can be. Many of our major choices in experimental design were made with the goal of increasing the repetition rate and I will point them out in Chapter 2.

1.2.4 Phase diagram and correlations in different regimes

Equation (1.44) is proportional to the likelihood of finding two particles in the same place. When finding the second one is statistically independent of finding the first, $g^{(2)}(z) = 1$, while $g^{(2)}(z) = 0$ indicates a certainty that no second particle will be found. Clearly this should tell us something about interactions. When the particles do not interact, particles can be found in the same location. If the phases are incoherent, as at higher temperatures, detecting a particle at $\boldsymbol{r}$ location does not diminish the chances of finding another particle at nearby $\boldsymbol{r}'$. In fact, it makes it more likely because the probability amplitudes for detecting each boson at the other location interfere constructively. This becomes even more true as $\boldsymbol{r}' \to \boldsymbol{r}$. This is the Hanbury Brown and Twiss effect [110, 111], also known as bunching. It is expected to be maximal in a thermal, ideal Bose gas where $g^{(2)}(z) = 2$. Interactions impose an energy cost to particles being near each other so finding a particle at $\boldsymbol{r}$ makes it less likely to bring another there. If the interactions become very strong, it becomes impossible and $g^{(2)}(z) \to 0$ if $\gamma \to \infty$ and the temperature is not so high as to overcome this interaction energy cost. Aside from these there is the *quasi-condensate regime*. Here, the temperature is low enough and the interactions are present, but weak enough, so that finding a particle at $\boldsymbol{r}$ does not give very much information about finding another one at $\boldsymbol{r}'$ and $g^{(2)}(z) \approx 1$.

The regions can be further divided by whether fluctuations from the interaction energy or kinetic energy is more prevalent. I will examine the case for uniform density. In the harmonic

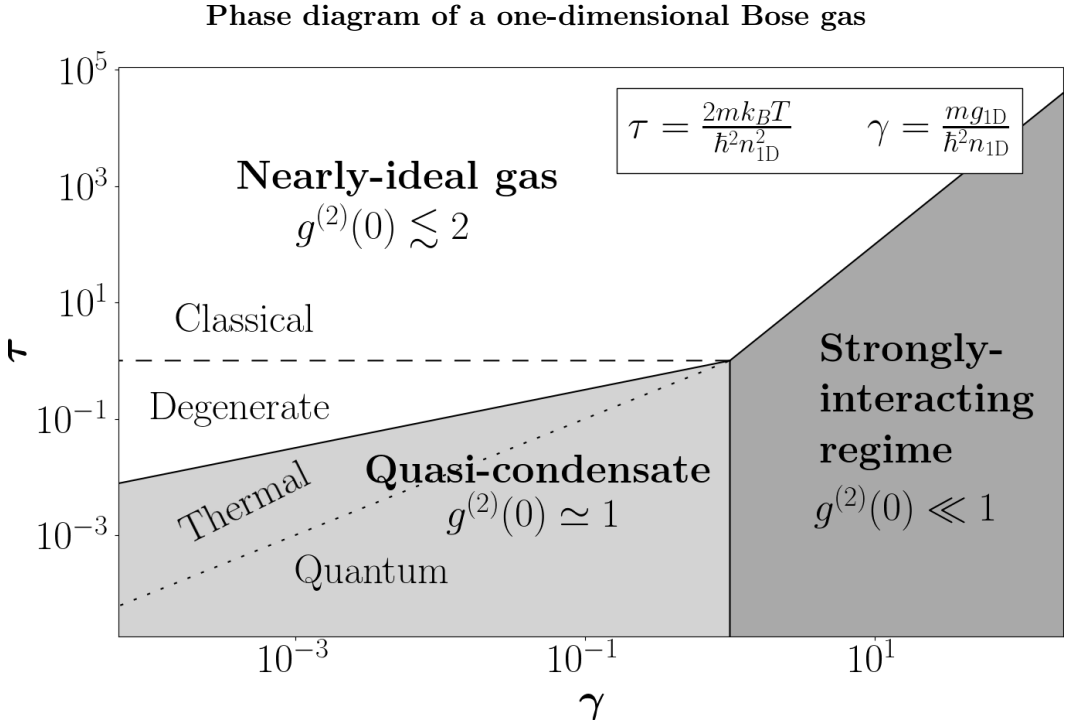

Figure 1.5 — One-dimensional phase diagram. Axes are defined in the inset from (1.34) and (1.45). Regimes are discussed in § 1.2.4. Recall that the boundaries between regions are not true phase transitions associated with broken symmetries and are instead smooth crossovers and boundary lines on the graph are meant to serve as guides.

trap with inhomogeneous density, different parts of the gas can be in different regimes but a local density approximation, $n \to n(z)$, can be made. To tell which kind of fluctuations dominate, it will be useful to define a temperature scale relative to the kinetic energy to go with γ (defined in equation (1.34)) which defines the interaction regime. To this end, we define the *reduced temperature*, τ, [56] as the thermal energy scaled by the kinetic energy scale, E_{kin}. This reads

$$\tau = \frac{k_B T}{E_{\text{kin}}} = \frac{2mk_B T}{\hbar^2 n_{1D}^2} \tag{1.45}$$

where $E_{\text{kin}} = \hbar^2 n_{1D}^2 / 2m$ as before. This should give a sense of the energy of thermal fluctuations in the gas without counting coherent excitations.

It will also be useful to have a length scale, l_c, on which we expect correlations to decay. This will be different depending on the regime which the gas is in and I will give values for the different regimes in the relevant sections.

Now we have everything we need to talk about the different regimes that the gas can be in and how the correlation functions behave in them. At finite temperature there are three overall regimes labeled in bold text in Figure 1.5, two of which have sub-domains related to what kind fluctuations are giving rise to the value of $g^{(2)}(z)$ which characterizes the region. I will discuss the full phase diagram at finite temperature, touching upon the transition to the classical gas, and then discuss a few features of the zero temperature case since apart from these features it is just a sub-domain of the finite-temperature case. I will use the $g^{(2)}(z = 0)$, at the point of maximum density to characterize the regions, but the reader should be aware

that in an inhomogeneous gas could have different parts in different phase regimes. Details are in [56].

Strongly interacting regime This is also called the *strong coupling* or *Tonks-Girardeau* regime when it is in one dimension after the physicists who first characterized one dimensional gases of hard spheres in the classical [112] and quantum [86] cases respectively. This is where $\gamma \gtrsim 1$ or $\sqrt{\tau}$, whichever is larger. Here, the energy cost imposed by interactions is high, so particles tend to spread out from each other, reducing correlations and $g^{(2)}(0) \to 0$ for $\gamma \to \infty$. The length scale $l_c \sim 1/n_{1\mathrm{D}}$ because the particles should be evenly spaced. It should be noted that particles can be spread out both by high interaction energy *or* by low density. This is captured by γ (1.34) depending on the density. In this regime, it is possible to make a map to fermions [86] since it is so unlikely to find two particles at the same position, it functions like the Pauli exclusion principle. It is in this region where γ is high that the Lieb II branch excitations are expected to be the most prominent, taking up more of the kinetic energy for a given momentum as γ is increased [49, Fig. 3].

Systems of strongly interacting bosons have been made in arrays of tubes [51, 53] with lattices along the long axis that lead to quasi-particles moving in one dimension with strong *effective* interactions by increasing the *effective* mass [113]; or in tubes where the longitudinal confinement is tuned to change γ, which depends on the one-dimensional density [52, 114]. Currently, it is difficult to go deep into the strongly interacting regime with atom chip experiments. The crossover has been probed on an atom chip [54]. That experiment relied on special techniques to increase the transverse trapping frequency [55]. They decrease $a_\perp$ and drive up the interaction strength (1.30). However these techniques cannot drive it arbitrarily high and have not produced systems deep in the strongly interacting regime. The hope is that the experiment presented in the thesis will provide a way to directly probe a gas of a single species of boson well in the strongly interacting regime.

Quasi-condensate regime This is also called the *intermediate coupling*, *weakly interacting*, or *Gross–Pitaevskii* regime since it is the regime where the Gross–Pitaevskii equation (1.25) is valid. It is characterized by more than just weak interactions, since the ideal Bose gas would also have weak interactions (i.e. *none* in the strict sense of "ideal"). This regime also requires that the temperature be low such that $\tau^2 \lesssim \gamma \lesssim 1$. In this regime, changes can be expected on the order of the healing length, $l_c = \xi$ given in equation (1.22), since that is what it was defined for. As discussed above, the phase fluctuates and the phase correlations $g^{(1)}(0)$ fall off as in equation (1.40). How the phase coherence length is determined depends on whether the fluctuations that cause it to fall off are quantum or thermal fluctuations. In the *quantum degenerate* regime, where $\tau \ll \gamma \lesssim 1$, quantum fluctuations suppress the phase coherence and [57]

$$\ell_\phi \sim \xi \exp(\pi/\sqrt{\gamma}). \tag{1.46}$$

When $\tau \gg \gamma \gg \tau^2$ the phase coherence is destroyed by thermal fluctuations, the gas is in the *thermal degenerate* regime or is a *thermal quasi-condensate*, and $\ell_\phi = \hbar^2 n_{1\mathrm{D}}/mk_BT$ [115]. In both cases, $\ell_\phi \gg l_c = \xi$ and there will be some noticeable (if not long-range) phase coherence. This will make $g^{(2)}(0) \approx 1$ with different corrections depending on whether the decoherence is caused by thermal or quantum fluctuations. In the quantum case, the correlation function will be

$$g^{(2)}(0) = 1 - 2\frac{\sqrt{\gamma}}{\pi} + \frac{\pi\tau^2}{24\gamma^{3/4}} \quad \text{for} \quad \tau \ll \gamma \ll 1. \tag{1.47}$$

While in the thermal condensate regime

$$g^{(2)}(0) = 1 + \frac{\tau}{2\sqrt{\gamma}} \quad \text{for} \quad \gamma \ll \tau \ll \sqrt{\gamma}. \tag{1.48}$$

All of the trailing terms are small.

This regime, with both weak interactions and quantum degeneracy has generated considerable interest. Because the gas is degenerate, it still has quantum features like short-scale phase coherence and long-range density correlations. Because the interactions are weak, mean-field theory applies and the Gross–Pitaevskii equation (1.25) can be used This makes this an attractive regime for experiments to search for phase coherence [92, 116–119], test the excitation spectrum predicted by Bogoliubov theory [120], look for excitation phenomena like solitons [121, 122] and study superfluidity [123, 124].

Decoherent regime This is also called the *nearly-ideal gas* regime because fluctuations of both density and phase have taken over coherence of either kind and the gas behaves almost like an ideal, non-interacting, gas. This happens when $\gamma \lesssim \tau^2$ or $\sqrt{\tau}$ whichever is lower. If $\tau < 1$ the gas is still in the *decoherent quantum* regime, where kinetic excitations dominate over thermal fluctuations, while when $\tau > 1$ the gas is in the *decoherent classical* regime. Wick's theorem [125] can be used to show that the correlation function of free, non-interacting bosons should be $g^{(2)}(0) = 2$. The classical or quantum subdomains of the decoherent regime correspond to different corrections to this figure. In the decoherent quantum regime

$$g^{(2)}(0) = 2 - \frac{4\gamma}{\tau^2} \quad \text{for} \quad \sqrt{\gamma} \ll \tau \ll 1 \tag{1.49}$$

while in the classical regime

$$g^{(2)}(0) = 2 - \gamma\sqrt{\frac{2\pi}{\tau}} \quad \text{for} \quad \tau \gg 1, \gamma^2. \tag{1.50}$$

In the decoherent classical regime, the de Broglie wavelength (1.13) becomes smaller than the one-dimensional scattering length (1.29) and so quantum interaction effects disappear and the gas can be treated classically.

Zero temperature transitions At zero temperature, the gas will be degenerate, that is, in one of the two coherent regimes listed above, with one crossover between the strong and weakly interacting regimes around $\gamma = 1$. At zero temperature there is not enough energy to thermally excite any particles and at zero interaction there can be no quantum depletion so all of the particles are in the ground state even if the excited states are in-principle thermodynamically accessible. Here, one gets something like a true BEC, even in the thermodynamic limit since every single particle is in the ground state. It is not precisely the same however since it does not come from the excited states being saturated. In the weakly interacting regime, there is phase coherence which falls off as in equation (1.46), which means that the phase coherence gets longer as interactions, γ, get less significant. For a system of finite size, it is possible to have a phase coherence length on order of the system size $\ell_\phi \gtrsim L$ and therefore have something that acts like a true condensate in practice, even if that description is not valid in the thermodynamic limit [126].

None of the boundaries between these regions are associated with a phase transition except the transition away from $\gamma, T = 0$ [56], and even this is arguable because there is no symmetry

broken when they are crossed and so the regimes can blend smoothly into each other.[6] To experimentally determine which regime the gas is in, $g^{(2)}(0)$ must be calculated by gathering many data points over which to average, as discussed previously, and its behavior matched to the relations above. To do that, and to move between different phase regions, one needs control over both the temperature and the interaction parameter γ. Although γ can be affected by the density at a fixed interaction strength (that is the approach taken in that past in the experiments cited above) we hope to be able to control interaction strength g_{3D} directly, in a manner I will discuss in the next section. The cooling techniques we will use to control the temperature and shift τ are fairly standard and I will touch on them in the next chapter. The technique we will use to control the interaction strength, g_{3D} and thereby γ at all densities is what I will describe next.

1.3 Controlling interaction strength: microwave induced Fano–Feshbach resonance

Fano–Feshbach resonance (FFR) or just *Feshbach resonance* is a term used to cover phenomena resulting from the state of elastic collision between a pair of atoms being coupled in one way or another to their bound molecular state. The *location* of the resonance refers to the point where this coupling is strongest. The basic conception in cold-atom physics is that the energy of a colliding state and a bound state must be brought very close together. This is usually done by adjusting a background magnetic field which will affect the potential felt by atoms in different spin states differently [37, 38, 127]. It can also be done by providing photons which couple the two states. In general this is called either an *optical* or *photo-induced Feshbach resonance* (especially if the photons are in the visible wavelengths) or *photoassociation* (especially when it us used to produce molecules instead of just tune the interaction strength). See the review by K. Jones et al. [42] for more details.

The novel Feshbach resonance we intend to use can be considered a kind of light-induced Feshbach resonance but with the photons in the microwave range. It was proposed in 2010 by Papoular et al. [43] who calculated the location and width of the resonance for the alkali metals up cesium. This is why we have included a microwave guide on the atom chip.

In § 1.3.1, I will be covering the bare minimum of background about cold collisions; just enough to give some sense of what is being coupled to what. For these results I will be relying on a course by J. Dalibard [128] with supplemental references where appropriate. Then, in § 1.3.2, I will talk about the usual magnetic Fano–Feshbach resonance widely used in cold atom experiments, drawing on two reviews, one by C. Chin [32] and another by T. Kohler [129]. Then in the last part, § 1.3.3, I will deal with the specific case of microwave-induced Feshbach resonance on which there is naturally very little literature apart from the paper by Papoular et al. [43]. Implications for the experiment will be discussed at the end.

1.3.1 Basics of cold collisions

A cold collision is a quantum scattering problem (that is the collision part) at very low kinetic energy (the "cold" part). For the basic framework of quantum scattering theory, see [71, Ch. XVII] but the core of the method is to transform the problem of two particles

6. A symmetry cannot break smoothly as there is no sense in speaking of a "partially broken symmetry." A symmetry is either present or absent and if absent, whatever was accounted for by the symmetry must be kept track of with the order parameter, no matter how small.

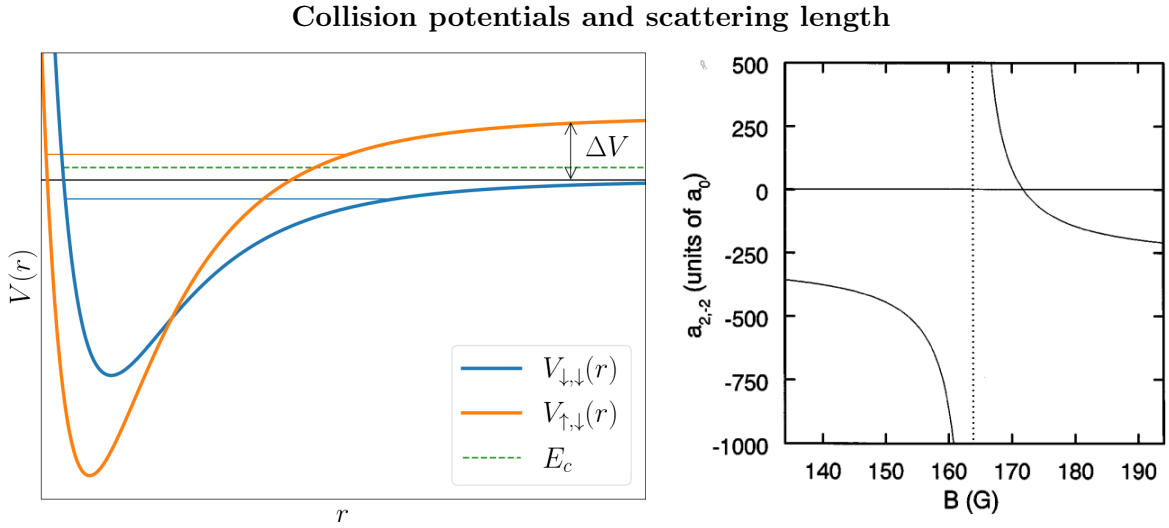

Figure 1.6 — Left – Example inter-atomic potentials for atoms colliding at energy E_c. The shallowest bound states are marked with horizontal lines. Right – Scattering length $a_{2,-2}$, for ^{85}Rb atoms with one in the $|F = 2, m_F = +2\rangle$ state and, the other in the $|F = 2, m_F = -2\rangle$ state, given in terms of the scattering length in the open channel, a_0, where the open channel is when both atoms are in the $|F = 2, m_F = -2\rangle$ state. The location of the Feshbach resonance where the scattering length diverges is highlighted with a dashed line. Adapted from [39].

of mass m_1, m_2 scattering off each other into a problem of a single particle of mass $\tilde{m} = m_1 m_2/(m_1 + m_2)$ scattering off of a central potential, $V(\boldsymbol{r})$. Here, $\boldsymbol{r}$ is measured from the center of mass of the two particles which is taken to be at rest. If the particles have no angular momentum relative to each other, as will be the case when the relative momentum, $k \to 0$, the potential $V(\boldsymbol{r})$ will be radially symmetric and, if the collisions are elastic, the wave function of the scattered states will be spherical harmonics with a phase shift relative to the wave function of the incoming state. This phase shift will be the one imparted by scattering off a hard sphere of radius a, the scattering length. That does not mean the potential itself is a hard sphere. The picture is modified somewhat in low dimension since the trapping potential is no longer symmetric. This means the angular momentum operator does no commute with the Hamiltonian. This leads to some corrections due to coupling between states with the same parity in their angular momentum [130]. At low relative momenta, these do not change the overall picture radically.

One can then construct a one dimensional radial potential after separation of variables. Let this be $V(r)$, which will be the inter-atomic potential and depend on the internal atomic states, resembling those depicted in Figure 1.6 and will depend on the internal state of the atoms. For the sake of example let us say the atoms can be in one of two states, $|\uparrow\rangle$ and $|\downarrow\rangle$, which can be thought of as spin states. The different combinations of these states lead to different shapes for $V(r)$. When the atomic states, for a given relative kinetic energy, are such that $r \to \infty$ is possible, this is called an *open channel*. An example of such a state with energy E_c is marked on the left of Figure 1.6. Combinations of states where the atoms would be bound for a given kinetic energy are called *closed channels*. An example would be the $V_{\uparrow,\downarrow}(r)$ potential in the same figure. Atoms with relative kinetic energy E_c would be bound by this potential if they were in the $|\uparrow,\downarrow\rangle$ state. A Feshbach resonance consists of narrowing

the energy gap between the open channel and a bound state of the closed channel. This is typically done with static magnetic fields which change the energy of the atomic spin states.

This is the usual picture when dealing with a Feshbach resonance but the difference between the open channel and the closed one can be one of symmetry[7] and "narrowing the energy gap" may be accomplished by coupling with a photon.[8] All that is required is that the open channel be strongly coupled to a bound state of the closed channel.

1.3.2 Fano–Feshbach resonance

The fact that the open and the closed channels rely on a difference in spin states means that they can be controlled via magnetic field using the *Zeeman effect*, the shifting of the energy levels of spin states by magnetic field due to those states having different magnetic moments. Then the gap between the asymptotic values of the potentials, ΔV, can be adjusted. For a Feshbach resonance to occur, it must be a little more than the bound states in the closed channel become energetically accessible at the collision energy, E_c. The energies must be *matched* not just that E_c is large enough. When this happens, the colliding state overlaps with a bound state and the particles close enough to interact are in a superposition with atoms bound in a molecule. If you think about what kind of radius a hard sphere would need to give the appropriate phase shift for a "scattering" problem that results in a molecule, you will see that at the resonance, the scattering length diverges.

For cold, elastic collisions, the dependence of the scattering length takes on the form [38]

$$a(B) = a_0 \left(1 + \frac{\Delta B}{B - B_0}\right) \tag{1.51}$$

where a_0 is the scattering length far from resonance, ΔB is the *resonance width*, and B_0 is the *resonance position* where a diverges. The width, ΔB, can take a negative value and this depends on which side of the resonance a diverges positively and on which side it diverges negatively. This results in relationships between scattering length and magnetic field that take the form of the right hand side of Figure 1.6. At the resonance, the scattering length diverges, to $+\infty$ on one side and to $-\infty$ on the other, and smoothly goes to the background value, a_0 far from the resonance, the speed of that return being given by the resonance width.

1.3.3 Microwave induced Fano–Feshbach resonance

The idea behind microwave induced Fano–Feshbach resonance (mw-FFR) is similar to photoassociation: the closed and open channel are not coupled by overlapping their wave functions but by providing photons that make of the right polarization to virtually populate the closed channel. The idea is illustrated on the left of Figure 1.7. This leads to the same kind of superposition that causes the scattering length to diverge in the ordinary Feshbach resonance. Recalling that this kind of resonance has been predicted for alkali atoms and taking a two-channel model where the spin $S = -1/2$ for both atoms in the singlet state in the open channel and one of the atoms having $S = +1/2$ in the closed channel, a formula similar to equation (1.51) can be derived. It is in terms of the frequency of the photons provided and it reads

$$a(\omega) = a_0 \left(1 + \frac{\Delta\omega}{\omega - \omega_0}\right), \tag{1.52}$$

7. In that case the resonance is called a *shape resonance* [32].
8. This is often called *photoassociation* [42]

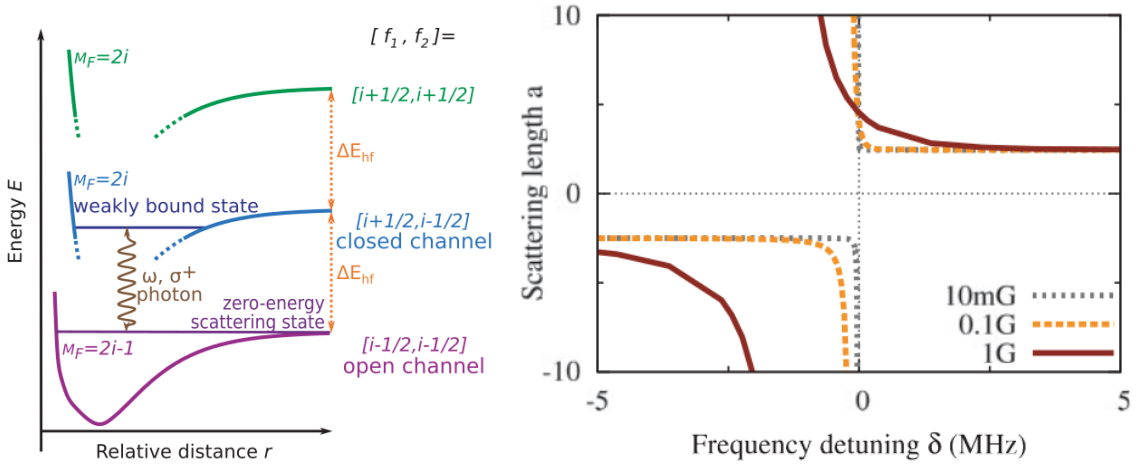

Figure 1.7 — Left – Potentials for the different collision channels and the microwave photon of frequency ω and polarization σ coupling the colliding state to the bound state of the closed channel. Right – Scattering length a, as a function of the detuning from the resonance, plotted at different values for the magnitude of the microwave field, which tends to broaden the resonance. Both from [43].

where the width of the resonance is now $\Delta\omega$ and its position is ω_0. This is plotted on the right in Figure 1.7 at different strength of the microwave magnetic field amplitude, B, which affects the width of the resonance. A full-channel model, where coupling between all states is allowed, results in only small corrections to this, as long as the kinetic energy in the collision is low enough.

Table 1.1 gives the main results from [43]. E_b is the energy of shallowest bound level in the closed channel. The splitting of the channel potentials is still subject to Zeeman effects, so the exact behavior so the resonance position is at frequency $\hbar\omega_0 \approx \Delta E_{HF} - |E_b| + \alpha B_0^2$ where ΔE_{HF} is the energy difference between the $S = -1/2$ and the $S = +1/2$ sublevels and the factor α is related to the coupling between the channels in the absence of photons. Lastly, $\Delta\omega$ is the resonance width which will scale with the amplitude of the microwave magnetic field like B^2.

Looking at Table 1.1 will help us find an alkali that will let us work with this resonance. Cesium has the largest width but it is difficult to condense, in fact impossible in a purely magnetic trap [22, 131]. The negative resonance width is clue to the fact that at zero magnetic field, the scattering length is negative [132], part of the difficulty in forming a stable condensate. The next broadest is sodium, which presents no such difficulties,[9] motivated our choice to work with it.

1.4 Conclusion

This chapter makes clear the kind of system we are going to study, the physics we hope to test, and give an idea of the ways we have to test them. We are going to make a one-dimensional system by confining sodium atoms in an elongated harmonic trap, with very tight confinement

9. The first sodium BEC [133] used a magnetic trap with an optical "plug" to prevent spin-flip losses. We prevent these by one trapping one spin state, see § 2.2.4.

Microwave induced FFR critical values

	^{7}Li	^{23}Na	^{41}K	^{87}Rb	^{133}Cs
$\lvert E_b\rvert/h$ (MHz)	12 000	200	140	25	110
$\omega_0/2\pi$ (GHz)	12	1.6	0.12	6.8	9.1
α (kHz $\cdot$ G^{-2})	0.33	6.8	21	120	30
$\Delta\omega/2\pi$ (Hz)	6	1400	350	60	−4500

Table 1.1 — Reproduced from [43]. Quantities explained in the text. The resonance width $\Delta\omega$ is given for when the microwave magnetic field amplitude is 1 G

in two dimensions, and then cooling the system enough, so there will not be enough energy to excite any atoms in the transverse, tightly confined direction. In this configuration, we will not be able to achieve true Bose–Einstein condensation; there will always be particles in excited states in the long direction since we will not be able to achieve a system with no interaction and thermal energy. This means that phase coherence, measured by $g^{(1)}(0)$ (1.37), will fall off exponentially. Still, there are some regimes, the quasi-condensate regime, where the interactions are weak enough and the temperature low enough that the phase coherence can be significant. The different regions of the phase diagram are closely associated with different behavior of the density correlations measured by $g^{(2)}(0)$ which is around one in the quasi-condensate regime; close to two when the gas is an ideal Bose gas, taking the value two when the gas is perfectly ideal and non-interacting; and falls to zero in the strongly-interacting Tonks-Girardeau regime, which has, so far, not been deeply probed directly using a single species, but either in lattices [53] or just on the edge of the regime [94]. A single tube of a single species, far in the strongly-interacting regime has yet to be made to compare with the exact Lieb–Liniger solution. A yet closer comparison would be to make a one-dimensional box potential. Then the local density approximation and generalized hydrodynamics can be done away with and exact comparisons can be made. The correlation functions used to distinguish the different regimes in the gas are defined via ensemble averages, so we will need to take data over many instances of the system. This demands a high repetition rate, a consideration which will drive many of our experimental choices.

To explore these different phase regions we will need to control the temperature and the interaction strength, and measure the correlation functions. We will control the temperature with the standard technique of evaporative cooling that I will say more about in the next chapter. The interaction strength (1.34) we hope to control by directly tuning the scattering length, as opposed to manipulating the density, via a predicted but yet to be confirmed method: microwave induced Fano–Feshbach resonance, which motivated our choice to work with sodium and on an atom chip. Atom chips trap atoms with magnetic fields close to their surface. The magnetic fields needed to produce a Feshbach resonance in sodium require magnetic fields over 100 G [39] or even almost 1000 G [127]. This can make it difficult for the fields generated by the wires to compete and can make trapping difficult.

Chapter 2

Experimental overview

When I arrived, the experimental setup had been completed and put under vacuum up to the valve (marked V in Figure 2.2) in the transport chain (F). This means the laser system (A and B) was operational, the oven (C) and initial cooling stage (D and E) was operational, and the magnetic transport could carry atoms from the initial trapping chamber (E, see page 51) to the valve (V) and back where they were imaged, albeit with a camera body intended for other purposes (see § 2.3.1.3). What remained to be done was the cleaning, assembly and installation of the science chamber (G); the manufacture, and assembly of the chip tower as well as the mounting and testing of the atom chip (see Chapter 3). There were and are other miscellaneous improvements I will point out where appropriate.

The purpose of this chapter is give a general overview of our experiment: orienting the reader with respect to the machine, giving critical values, justifying our design choices, and explaining what we will measure and what we hope that will tell us. For more technical details on the construction, the reader is referred to the thesis of D. Ben Ali [1], the previous doctoral student on the project, and to the paper [134] for particular treatment of the permanent-magnet Zeeman slower. First, in Section 2.1 I will connect to the last chapter by describing the pressures the physics we are trying to observe puts on our choices. Then in Section 2.2 I will go through each section of the setup, noting key values and noting a few difficulties and how they were resolved. Lastly in Section 2.3 I will describe some of the measurements we hope to make, how we will make them and what we hope they will tell us.

2.1 Experimental goals and considerations

In a sentence, we want to explore the phase diagram of a one dimensional system using the predicted but yet-to-be observed microwave-induced Fano-Feshbach resonance (see Section 1.3) to tune the interaction strength. Starting from this novel Feshbach resonance, one can see from Table 1.1 that the alkali that has both a broad resonance and is relatively simple to condense is sodium (see arguments against cesium in § 1.3.3). We chose to work on the $3^2S_{1/2} \rightarrow 3^2P_{3/2}$ transition, called the D_2 line. We chose this because of the presence of a cycling transition, $|F = 2\rangle \leftrightarrow |F' = 3\rangle$, absent in the D_1, $3^2S_{1/2} \rightarrow 3^2P_{1/2}$, transition [135]. The hyperfine structure of the chosen transition, along with the various frequencies we have our lasers tuned to, is depicted in Figure 2.1. I will come back to these laser frequencies in

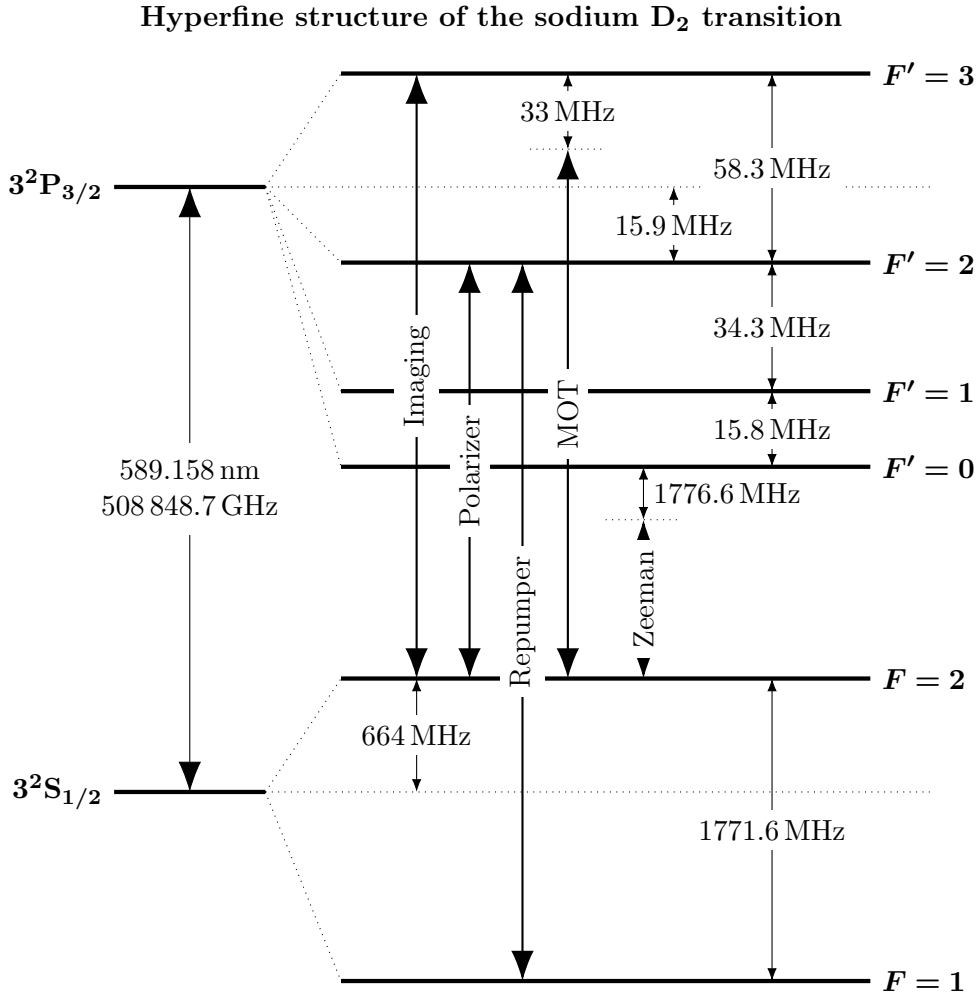

Figure 2.1 — Level structure of sodium with our laser beam frequencies superimposed. Adapted from [1, Fig. 1.1] and [135, Fig. 2].

§ 2.2.2 after I cover more general considerations.

We want to directly probe a system of bosons with one-dimensional dynamics. This means that we must make an elongated potential, tightly confined in two dimensions (see § 1.2.1). There are a two main ways to do this: arrays of tubes made by interfering optical beams [52, 136–139] or with an atom chip [44, 47, 92, 105, 120, 140, 141]; although there are proposals to form highly elongated traps with thin films of permanent magnets that rely on strong gradients near domain boundaries [142], which has the disadvantage of being impossible to switch off, or to make optical dipole traps with Fresnel lenses [143]. In the tube array experiments, the interaction parameter, γ, (1.34) is determined by averaging over all of the tubes in the array [144]. We want to directly probe a single system in the different phase regimes rather than make inferences from averages, and we hope that we will be able to tune the interaction strength itself with this novel Feshbach resonance rather than having to rely on changing the strength of the confining potential, thus manipulating $a_{1\mathrm{D}} = a_\perp/a$ by changing $a_\perp$ (1.29) [52]. This would suggest using an atom chip, which provides other advantages. Since atom chips are made of gold wires laid down on a substrate, we can incorporate coplanar waveguides (see 3.1.3) to deliver the microwaves as well as radio frequency wires that can dress

the trap. We are also not strictly limited to a single tube and have incorporated a second trapping wire so that we can make an elongated double-well potential in a manner similar to [145] or using radio frequency fields as in [146].

The other major consideration for us, which should be clear from the last chapter, is the repetition rate. We are trying to explore a phase diagram that will not have any true phase transitions and is characterized entirely by smooth crossovers between regions marked out by different values in the second order correlation function, $g^{(2)}$, (1.41). Since this is defined in terms of an ensemble average, many instances of the system must be measured, that is to say, the ensemble must be well-sampled.[1]

The small wires on the chip provide steep gradients but not very strong fields so the atoms must be loaded onto the chip while already cold. Methods for loading atom chips can be broken into two broad categories [45, §2.6]. One is the practice of loading the cold atoms from a magneto-optical trap (MOT) right on top of the chip. A typical MOT comprises six beams in three counter-propagating pairs, one for each axis. The ones made on chips can use the surface of the chip to reflect some of the beams (a *mirror MOT* [147]). Often these are used with some prior cooling stage like a two-dimensional MOT [148, 149] which only confines in two dimensions and thus makes a beam of cold atoms. The other method is to cool the atoms in a three-dimensional MOT in another chamber and transport the cold atoms to the chip (see references in the first paragraph of Chapter 5). We have chosen this latter option because it allows us to load a MOT for cooling at the same time as we are condensing and performing observations on the atoms trapped by the chip. This is one reason why the transport has a corner in it: to isolate the MOT chamber from the science chamber and prevent light scatted by the MOT from reaching the atoms in the science chamber and ejecting them from a magnetic trap in the chip chamber.

The next question is how to load the MOT in its own chamber. One way is to load that MOT from a vapor cell [149–152] which begins with atoms at room temperature but requires high density (and hence pressure) of the trapped atoms for efficient loading. This means that the lifetime in the MOT will be very low, or the pressure must be pulsed somehow. Attaining that high a pressure is also not easy for a sodium experiment since sodium has a low vapor pressure compared to other alkali metals.[2] As a rough comparison, K. Gibble et al. loaded their MOT from cesium vapor in 0.2 s [151] but cesium has around ten times the vapor pressure of sodium.[2] In order to raise the pressure we would have to heat the sodium, which would lower the capture rate which scales inversely with the most probable speed cubed [151, Eqn. 1]. This low vapor pressure, however, makes sodium very well-suited to putting into an oven and forming an atomic beam.

Then this beam can be slowed with a *Zeeman slower* [10], which uses the radiation pressure from a laser to slow the atoms and a magnetic field to take advantage of the Zeeman effect to keep the the atoms on-resonance as the Doppler shift they see in the light changes. While loading from a vapor can be fast (from 0.2 s [151] to a second or two [149] as compared to our loading times of around 10 s [1, Fig. 3.14]) we estimated that the technical challenge of getting a two-dimensional MOT or a vapor cell working with the low vapor pressure of sodium was not worth the effort when the time we will be loading, cooling, manipulating and imaging the atoms on the chip will already be on the order of seconds anyway.

This gives the overall shape of the machine: we need a laser source that addresses various

1. If we could show that the system is ergodic, the main concern of Part II, we would only need to observe the system for a long time, which would be a different set of challenges.
2. 89 mbar at 950 K [153] versus 437 mbar for K [154], 500 mbar at 925 K for Rb [155], and 908 mbar at 950 K for Cs [156]

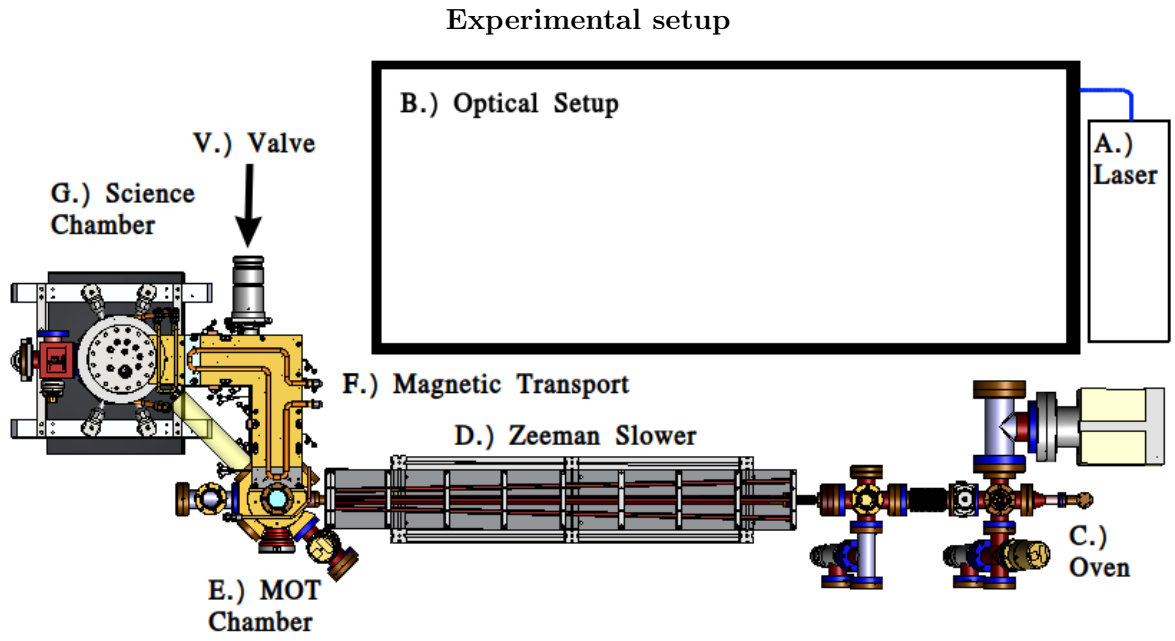

Figure 2.2 — Top-down schematic of the experimental setup. Optical fibers running from the optical setup box to the experiment omitted for clarity.

transitions in the sub-structure of the sodium D_2 line. We will feed a MOT with a Zeeman slower, and this MOT will be disconnected as much as possible from the science chamber with a method of transport between them that does not sacrifice the low temperature or number of atoms attained by the MOT or lose us too much time. The atomic cloud will need to be imaged at various stages of this process and the final experiments will be performed on an atom chip in the science chamber. Let us go through these in a little more detail.

2.2 Overview of the setup

Here I name and briefly define all of the parts of our experimental setup. After a brief note on experimental control, I follow the same structure as the labeling in Figure 2.2: starting with the laser (A) and optical setup (B), then following the path of the atoms from the oven (C) to the science chamber (G). For explanations of any of the optical elements, I recommend the excellent textbook [157]. I will make references to this where I can and include other references where appropriate.

2.2.1 Computer control

The experiment is programmed by a suite of MATLAB scripts written by A. Perrin which in turn programs an ADwin-Pro-II control system from Jäger Computergesteuerte Messtechnik GmbH. This system sends out analog signals and transistor–transistor logic (TTL) pulses according to a sequence that is loaded into the memory of the ADwin by MATLAB. The analod outputs handle things like the setting and ramping of frequencies for rf-drivers and voltages for current supplies and the digital TTL signals send signals to switch states. Crucially, they send the signal to switch between pairs of coils in the magnetic trap chain (e.g. the vertical lines in the plot to the right in 5.4). Once the command to execute the sequence is given,

the ADwin system sends out the sequence of signals with no further need for any input, thus eliminating any delay or error that could come from a the PC or the USB connection. For more details on the control interface see [1, §1.3.2].

2.2.2 Laser and optical setup

Our laser system [1, §1.1.2] is based on a tunable diode laser, from Toptica Photonics.[3] This makes a 1178 nm beam of around 50 mW which provides the "signal" into a single frequency visible *Raman fiber amplifier* system from MPB Communications.[4] It includes the Raman amplifier, the ytterbium fiber laser to "pump" it and a second harmonic generation (SHG) device. A Raman fiber amplifier consists of a fiber and a pump laser with a different wavelength from the signal [158, p. 755]. In the fiber, the photons undergo inelastic stimulated Raman scattering where pump photons are converted into photons at the signal frequency, thus boosting its power. The ytterbium fiber laser that pumps the amplifier comprises two laser diodes, one of wavelength 0.9 µm and one of 1.0 µm, and a fibre doped with Yb^{3+} ions. The light from the diodes excites the ions which emit at 1.12 µm which pumps the Raman amplifier. The light emerges from the amplifier maintaining a wavelength of 1178 nm but now having around 8 W of power. This is fed into a crystal and undergoes second harmonic generation [158, §1.2.1], wherein two photons of one frequency are converted into one photon of twice the frequency. The process is not perfect however and the beam emerges at $\lambda = 589$ nm but with less than 2 W of power. We typically use a beam of 1.27 W and, according to the manufacturer, 2 W is the maximum power we could expect but we should not try to run it above 1.8 W.

This is fed into the optical setup [1, §1.1.4] (B in Figure 2.2) via a polarization preserving[5] optical fiber. Upon exiting the second harmonic generation crystal and entering the optical setup, the laser is locked to a frequency 267 MHz below the $|F = 2\rangle \leftrightarrow |F' = 3\rangle$ cycling transition. I will follow D. Ben Ali's convention of giving frequencies in terms of their offset from the frequency of the cycling transition, ν_0. The beam is then drawn off into four separate stages where its frequency is conditioned and, if necessary, it is split up further into the different beams we need for the experiment. Polarizing beamsplitter cubes [157, §4.9.9] or just polarizing beamsplitters (PBSs), half-wave ($\lambda/2$) plates [157, §6.9.3] are used to control the splitting of the beam and various singlet lenses are used to keep the beam from spreading, but these do not necessitate further remark.

First, around 385 mW is siphoned off for the Zeeman beam. It passes through a cascade of four *acousto-optical modulators* [157, §7.10.4] (AOMs). These devices set up a diffraction grating of standing waves at a driving frequency, ν_A. The beam emerges from them split into different orders where the beam of the n^{th} order is frequency shifted by $n\nu_A$. Using a mirror in a cat's eye configuration we can make a *double-pass* AOM [159] which returns a beam back along the path it came, shifted by double the driving frequency. We send the beam twice through an AA Opto-electronic brand AOM.[6] This is the AOM we tune when we want to adjust the frequency of the Zeeman beam. We then send it through a cascade of three AOMs from Gooch & Housego (G&H),[7] passing twice through each. Two adjustable mirrors inject the beam into a polarization preserving optical fiber; as it is with all our fiber injections.

3. Model DL100/pro design.
4. Model VRFA-P-1800-589-SF.
5. all of our fibers are polarization preserving so I will drop the modifier.
6. Model MT200-A0.5-VIS.
7. Model 3200-120.

Adjustable $\lambda/2$ plates let us ensure that the polarization of the beam is matched to axis of of the fiber. If it is not, it can lead to poor stability of the polarization at the exit [157, §8.4.4]. This, in turn leads to poor stability in the intensity after the polarization necessary for the experiment is filtered for.

All four of these AOMs have central frequency of 200 MHz. Three are driven at around 205 MHz while the one used to tune the Zeeman frequency is currently driven at 193 MHz. The negative first order is taken so this shifts the frequency of the Zeeman beam down by just over 1.6 GHz, putting it almost 1.9 GHz below the cycling transition to be on resonance with the atoms counter-propagating to the light at $(870 \pm 3)\,\mathrm{m \cdot s^{-1}}$ [1, Tbl. 2.3]. This large detuning means that the Zeeman beam will not interact with the slower atoms in the MOT. After all of these AOMs, around 80.3 mW enters the fiber and just under 60 mW makes it into the beam. This long chain of AOMs is, in fact, cheaper and more practical than a single AOM at 1.6 GHz, which have to be custom-made, or running a 1.7 GHz AOM 100 MHz away from its central frequency. This problem still motivates instrument research [160].

Next 700 mW of the power for the imaging and MOT beams is taken. All of this is passed twice though an AA Opto-electronic AOM[8] tuned to 117 MHz and the first order is taken so that the beam comes out at 33 MHz below ν_0, just the detuning needed for our MOT [12] [1, Ch. 3]. The minimum amount allowed by our PBS and $\lambda/2$ scheme, just under 30 mW, is taken for the imaging beam. This is passed through a $\lambda/2$-plate, a neutral density filter which cuts the power by a factor of ten, a PBS, and then another $\lambda/2$-plate. This branch (like the others) is siphoned off of the main beam by being reflected off of a PBS. The polarization of the reflected beam from a PBS is not always very pure [157, §4.9.9] and this sandwich of $\lambda/2$ wave plates is made to stabilize the polarization. It does this by taking from a beam reflected by a PBS to a beam transmitted by one but keeping the same polarization via the $\lambda/2$-plates. The imaging beam must be very low power to avoid saturating pixels in the CCD sensor and we wind up with 254 µW.

All the rest of the power in this branch is passed through a QUBIG[9] electro-optic modulator [161, §18.1B] (EOM). An EOM is a device which uses a voltage across a crystal with non-linear optical properties to apply a phase delay in proportion to the applied voltage. If we vary the voltage sinusoidally with frequency ω_{mod} we will get a wave emerging from the crystal which has the form

$$E(t) = Ae^{i[\omega t + B\sin(\omega_{\mathrm{mod}} t)]} \tag{2.1}$$

where ω is the frequency of the incoming wave and A and B are the amplitudes of the incoming wave and the modulation respectively. Expanding this to first order [78, Eqn. 4.2.1] and using the definition of the sine function [78, Eqn. 4.3.1] gives

$$E(t) = A\left(e^{i\omega t} + \frac{B}{2}e^{i(\omega+\omega_{\mathrm{mod}})t} - \frac{B}{2}e^{i(\omega-\omega_{\mathrm{mod}})t}\right), \tag{2.2}$$

the original beam with two sidebands separated from the main beam by the modulation frequency, one *blue-shifted* (that is up in frequency, toward the blue) and the other *red-shifted*. This should hold since the amplitude of the modulation can be assumed to be much smaller than the amplitude of the field associated with the laser.

By driving the EOM at 1.711 GHz we can add a blue-shifted sideband (a red as well, but this will be off-resonance with all transitions) to the MOT beam that will pump atoms from the $|F = 1\rangle$ state that would otherwise be a *dark state*, i.e. one that did not interact with

8. Model MT110-B50A1-VIS.
9. Model EO-Na23-3K.

the light. This state is dark while others are not because it is well-separated from the nearest state, $|F = 2\rangle$ by thousands of MHz. By contrast, $|F' = 2\rangle$ is only 58.3 MHz from the $|F' = 3\rangle$ state addressed by the laser. This closeness in the energy levels in the excited state means that the probability of a photon resonant with the $|F = 2\rangle \rightarrow |F' = 3\rangle$ cycling transition putting the atom in one of the other excited sub-levels is significant [135, Tbl. 8]. From here they can fall into the $|F = 1\rangle$ dark state and need photons that address that sublevel in order to *pump* (or *repump*) them out to, for example the $|F' = 3\rangle$ state where they can fall back into the cycling transition. In the MOT this is accomplished by the sidebands induced by the EOM. This MOT beam is fed into a fiber leading to a Schäfter-Kirchhoff brand "Fiber Port Cluster 1 to 6 for 589 nm" which breaks it up into six beams in six fibers with the appropriate polarizations. Each of these beams ends up with just under 30 mW each. I will discuss the functioning of the MOT in a little more detail in § 2.2.3.

The next branch of the optical setup is for the dedicated repumping beams for which 160 mW is diverted. We use two: one to prepare the atoms before the Zeeman slower and a recently-installed one for the MOT chamber to help with imaging. This entire branch is passed twice through an AA Opto-electric brand[10] AOM tuned to 100 MHz. The first order is taken so the beam is on the $|F = 2\rangle \rightarrow |F' = 2\rangle$ transition. It is then split into the repumper/polarizer beam for the Zeeman slower and the dedicated repumping beam for the MOT. The repumper/polarizer beam then passes through another QUBIG[11] EOM which adds sidebands at ± 1.77 GHz. This beam is to prepare the atoms for the Zeeman slower. Since it has a magnetic field transverse to the path of the atoms, the atoms need to be prepared to keep them on the $|F = 2, m_F = -2\rangle \rightarrow |F' = 3, m_F = -3\rangle$ transition in order to be cooled efficiently (see [134] or [1, Ch. 2] for details). So they are driven into the lowest m_F state with σ^- photons on the $|F = 2\rangle \rightarrow |F' = 2\rangle$ and pumped out of the dark state by the sidebands on the $|F = 1\rangle \rightarrow |F' = 2\rangle$ transition. This arrives at the experiment with a little more than 15 mW of power. The dedicated repumping beam passes through a Brimrose[12] AOM, tuned to 1.77 GHz putting it on-resonance with the $|F = 1\rangle \rightarrow |F' = 2\rangle$ transition. This AOM has a very low efficiency. It is rated at 15% efficiency, but we only manage to achieve 8% after several weeks of work. The beam from this is then injected into a fiber that takes it to a beam expander, from whence it is reflected downward, into the MOT, delivering 2 mW.

The remaining power in the laser amounts to around 40 mW and is fed into the saturated absorption mechanism used to lock the laser [1, §1.1.3] using a technique from [162] using a transition in iodine pointed out by [163]. As mentioned before, sodium has a very low vapor pressure. This means in order to get a dense enough gas in a vapor cell to use saturated absorption, we would need to heat the cell. Instead, we use a method that relies on the fact that there is a transition in molecular iodine, $^{127}I_2$, very close to the cycling transition in sodium.[13] A particular sub-level of this transition[14] is only 467 MHz above the cycling transition in sodium. We counter propagate two beams through an iodine cell, one higher-power one to saturate the transition and a pair of lower powered probe beams, one that travels along the path of the pump beam and one that does not. These two probe beams are measured by ThorLabs [15] photodiodes and the difference between them taken using a simple, passive

10. Model MT110-B50A1-VIS.
11. Model EO-Na23-3K.
12. Model TEF-1700-350.
13. the P 38 (15-2) transition [163] from the second ro-vibrational sub-state of the molecular ground state, $X : 0_g^+ (^1\Sigma_g^+)$, to the fifteenth ro-vibrational sub-state of the molecular excited state [164]
14. Specifically, the a_1 sub-level by the notation of [164].
15. Model PDA36A.

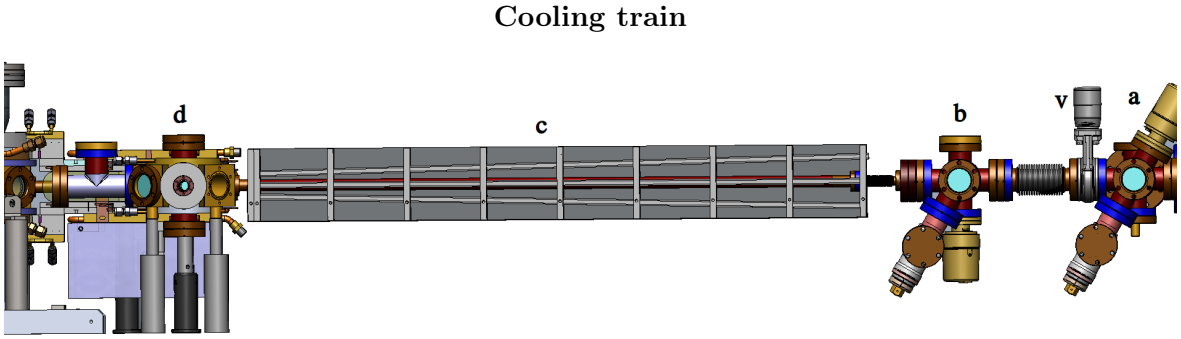

Figure 2.3 — Side view of the cooling train from the "cross" at the exit of the oven (a) to the MOT chamber (d). The repumper/polarizer beam enters at (b) going into the plane of the page, just before the permanent magnet Zeeman slower (c). A valve (v) keeps the oven isolated from the rest of the vacuum system when the atomic beam is not needed.

analog circuit. This difference provides the signal for the laser locking mechanism. By scanning the laser frequency using the piezoelectric crystal to adjust the cavity length and taking the difference between the probe beams, one obtains the spectrum of the iodine transition with peaks at the sub-levels where the beam that was traveling through the saturated iodine is brighter than the one which interacted with those transitions. The peak of the a_1 sub-level can then be used to lock the laser at this frequency with the help of a laser locking module from Toptica photonics.[16] For further details about how the frequencies are filtered, see [1, §1.1.3.4].

2.2.3 Cooling stage

Our oven is charged with pieces of metallic sodium. The procedure is described in [1, 1.2.4]. Absorption imaging can be done at the "cross" at the exit of the oven (marked "a" in Figure 2.3) to see if the oven needs refilling. A well-charged oven should produce a beam that absorbs $\approx 18\%$ of a low-intensity laser shined across it. The oven is operated at a temperature of 300 °C making an atomic beam with a velocity of $(870 \pm 3)\,\mathrm{m \cdot s^{-1}}$ [1, Tbl. 2.3]. When in operation, the pressure in the oven is on the order of 10^{-6} mbar. When not in use, it is cooled down to 120 °C and the pressure falls to $\sim 10^{-8}$ mbar. When shut completely off, the pressure can fall to that of the MOT chamber, on the order of 10^{-10} mbar or less. A valve (marked "v" in Figure 2.3) can be closed to isolate the oven for reloading.

Atoms leave the oven and are collimated into a beam by being forced through an opening 4 mm wide and then one 8 mm wide. This also lets us maintain a differential pressure between the oven and the rest of the vacuum system. The atoms then enter the Zeeman slower. This slows them down from an average velocity of almost $900\,\mathrm{m \cdot s^{-1}}$ to just under $30\,\mathrm{m \cdot s^{-1}}$ [1, Tbl. 2.3], slow enough to be captured by our MOT [12]. The MOT then further cools the atoms to $\lessapprox 200\,\mu\mathrm{K}$ (see §5.2.1 for temperature measurements); the atoms, trapped in the MOT, now having a well-defined temperature.The same coils that produce the magnetic field for the MOT are also used to form a magnetic trap, which forms the first step in the transport from the MOT chamber to the science chamber. The physics of Zeeman slowers [10] [75, §6.2.2] and magneto-optical traps [12] [75, §11.4] are well-understood, so I will not cover them in depth but rather point out a few particulars of our experiment so no large features go unexplained.

16. Model Digilock 110.

Zeeman slower Our Zeeman slower uses an arrangement of permanent bar magnets to generate a magnetic field that is transverse to the path of the atoms and increases in magnitude as they go. The fact that it is built out of permanent magnets means that there is no electrical power or coolant that needs to be supplied to it. However, because of the transverse magnetic field, the laser cannot be purely circularly polarized, which would excite atoms to a particular Zeeman level. This means that the laser can de-pump atoms into the dark state, at least until they are well inside the magnetic field of the Zeeman slower and the Zeeman splitting between the spin projection states becomes large. We minimize this by pumping them into the $|F = 2, m_F = -1\rangle \leftrightarrow |F' = 3, m_F = -3\rangle$ transition which is not coupled to the $|F = 1\rangle$ dark state. This is thanks to the selection rules [165, §2.2.1], which state that the electron angular momentum, $\boldsymbol{F}$, cannot change by more than one in a single jump. They are put on this transition by means of the repumping/polarizing beam mentioned above. This crosses the atomic beam at the window marked "b" in Figure 2.3, 15 cm before the start of the Zeeman slower, and takes the population of atoms that enter the Zeeman slower in the $|F = 2, m_F = -2\rangle$ state which will be efficiently cooled, from 3.7% to 31% [134]. This could be improved if the goemetry of the setup allowed for moving the repumper closer to the entrance of the Zeeman slower, giving the atoms less time to get de-pumped. Overall, the presence of the repumper/polarizer adds a factor of between three and four to the number of atoms trapped in the MOT. The fact that it is not a factor of eight is due to the fact that atoms in $m_F \neq -2$ states will still be slowed, they just have a higher probability of falling into the dark state. For further design particulars of our permanent magnet Zeeman slower, see [134].

MOT Chamber Our MOT is the typical one [12] using six counter-propagating beams, one pair for each axis, with circular polarization of opposite orientations. These beams are slightly red-detuned 20-30 MHz from the cycling transition with sidebands, one of which extends to the $|F = 1\rangle$ sublevel to pump atoms out of the dark state. This bath of photons forms an *optical molasses* [11] [75, Ch. 7] which provides a friction-like force on the atoms, slowing their motion. Trapping atoms by light fields alone would violate Earnshaw's theorem [166, 167] but a *quadrupole* magnetic field, one with a linear gradient and a magnetic zero at the center, generated by a pair of wire coils circumvents this, modifying the effective detuning of the atoms from the light to add a restoring force that traps them. When we are ready to transfer the atoms into a purely magnetic trap, we ramp up the current in the coil creating the quadrupole field from 5 A to ~ 70 A taking the gradient of the magnetic field in the horizontal plane from around $5\,\mathrm{G \cdot cm^{-1}}$ to $65\,\mathrm{G \cdot cm^{-1}}$. The pure magnetic trap will only hold certain spin states (see § 5.1.1) so we drive the atoms into the $|F = 1\rangle$ dark state by shutting off the repumping sidebands. The decay rate from the excited state is $61.5 \times 10^6\,\mathrm{s^{-1}}$ [135, Tbl. 3] so it only takes a few micro seconds to drive the vast majority of atoms to the dark state. When the light is shut off, only the $m_F = -1$ sub-state is trapped (again, see § 5.1.1) so two-thirds of the atoms are lost. Now in the magnetic trap (this trap in the MOT chamber is what I will refer to as "the magnetic trap" unless otherwise noted) the atoms can be held, imaged, or moved into the magnetic transport that is the topic of Chapter 5. We can capture $\approx 6 \times 10^8$ atoms in around 10 s with this MOT and the oven at 300 °C.

2.2.4 Quadrupole magnetic trap

The magnetic trap works because an atom with a magnetic moment $\boldsymbol{\mu}$ under the influence of a magnetic field of $\boldsymbol{B}(\boldsymbol{r})$ at position $\boldsymbol{r}$ will have an interaction Hamiltonian written [165]

$$\hat{H}_B = -\hat{\boldsymbol{\mu}} \cdot \boldsymbol{B}(\boldsymbol{r}) \tag{2.3}$$

where $\hat{\boldsymbol{\mu}}$ is the operator associated with $\boldsymbol{\mu}$. If the magnetic field is so weak, and the magnetic moment sufficiently small, that the energy associated with $\hat{H}_B$ is small compared to the hyperfine splitting, then the energy levels of the atom can be described by $\hat{\boldsymbol{F}} = \hat{\boldsymbol{I}} + \hat{\boldsymbol{J}}$, the total angular momentum of the atom (the sum of the nuclear spin, $\hat{\boldsymbol{I}}$, and the electron's total angular momentum, $\hat{\boldsymbol{J}}$), and m_F the projection of $\boldsymbol{F}$ along the magnetic field which provides a quantization axis.[17] In this basis of hyperfine states, $\{|F, m_F\rangle\}$, the magnetic moment of the atom reads

$$\hat{\boldsymbol{\mu}} = -\frac{g_F \mu_B}{\hbar} \hat{\boldsymbol{F}} \tag{2.4}$$

where g_F is the Landé g-factor for the total angular momentum, μ_B is the Bohr magneton. As long as the magnetic moment of the atom adiabatically follows the magnetic field[18] the dot product in the interaction Hamiltonian (2.3) will project $\boldsymbol{F}$ onto $\boldsymbol{B}$ and yield a factor of $\hbar m_F B(\boldsymbol{r})$ and the potential felt by the atom at position $\boldsymbol{r}$ will be

$$V_B(\boldsymbol{r}) = g_F \mu_B m_F B(\boldsymbol{r}). \tag{2.5}$$

Atoms can then have internal states that are either low-field seeking (whose potential energy is proportional to the strength of the magnetic field, i.e. $g_F m_F > 0$) or high-field seeking (it is negatively proportional, i.e. $g_F m_F < 0$) [168]. Traps can then be formed by creating regions of space with a field minimum (to trap low-field seeking atoms) or a field maximum (to trap high-field seeking ones).

According to Wing's theorem, no region of space can have a single field maximum unless it contains a net charge [168]. Therefore, trapping high-field seeking atoms in our MOT chamber would require a magnetic field maximum inside the chamber. There are no magnetic charges, so we constructed a trap with a magnetic minimum and use atoms in low-field seeking states. In the ground state of sodium there are three such states [135]

- $|F = 1, m_F = -1\rangle$, $g_F m_F = 1/2$
- $|F = 2, m_F = +1\rangle$, $g_F m_F = 1/2$
- $|F = 2, m_F = +2\rangle$, $g_F m_F = 1$.

If we were to prepare the atoms in $|F = 2\rangle$, we would have to polarize the atoms in the $|F = 2, m_F = +2\rangle$ state to avoid a mixture of $m_F = 1$ and $m_F = 2$ states in the trap. Because, when these undergo collisions, it is possible for angular momentum to be exchanged between them [169] and the spin to switch to an untrapped state, incurring losses. To avoid these losses we prepare the atoms in the $|F = 1\rangle$ state before transferring to the magnetic trap.

To accomplish this, we turn off the EOM making the sideband on the MOT lasers which address the $|F = 1\rangle \rightarrow |F' = 2\rangle$ transition. Since $|F' = 3\rangle$ and $|F' = 2\rangle$ are only 60 MHz

17. If the magnetic field were stronger, it would force the electron spin, $\hat{\boldsymbol{S}}$, to be conserved separately from the electron orbital angular momentum, $\hat{\vec{L}}$, and we could not use only $\hat{\boldsymbol{J}} = \hat{\boldsymbol{S}} + \hat{\boldsymbol{L}}$.
18. If it does not, this can result in a Majorana spin flip which I will discus beginning on p4age 105

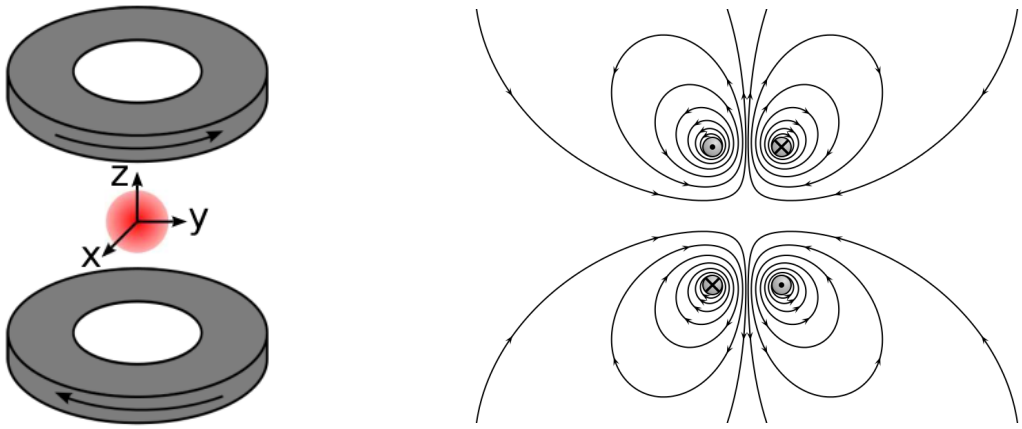

Figure 2.4 — Left — Coils in the anti-Helmholtz configuration; from [1]. Right — The magnetic field generated by such a pair.

apart, the main beam which addresses the $|F = 2\rangle \rightarrow |F' = 3\rangle$ cycling transition also puts some atoms into the $|F' = 2\rangle$ state. From there they fall into the $|F = 1\rangle$ state, which is a dark state now that the repumping sideband is turned off. In $\sim 10\,\mu s$ all of the atoms have fallen into the $|F = 1\rangle$ state with their F projections randomized among $m_F = 0, \pm 1$. Since $|F = 1, m_F = -1\rangle$ is the only trapped state left out of those remaining after this depumping procedure, only atoms in that state are left in the magnetic trap and 2/3 of the atoms in the MOT are lost when the magnetic trap is loaded, but the lifetime in the magnetic trap is extended.

Our trapping field takes the form of a quadrupole magnetic field generated by a pair of coils in an anti-Helmholtz configuration and is illustrated in Figure 2.4. An anti-Helmholtz coil, or coil pair, is pair of wire loops or coils arranged parallel to each other with their centers on an axis perpendicular to them, their centers separated such that the field between them is as even as possible,[19] connected in series so that equal amounts of current flows through them in opposite directions. This last condition is what distinguishes it from the "Helmholtz" configuration, where the current flows the same direction and the field more closely resembles that of a solenoid. Each of the loops creates a magnetic flux through its center outward from the space between them. Since the Maxwell-Thompson law requires that $\boldsymbol{\nabla} \cdot \boldsymbol{B} = 0$, this flux must enter the region in between the coils from the sides and, along with symmetry considerations, that a point where $\boldsymbol{B} = 0$ must lie at the geometric center of the coil pair, taken as the origin. This is the low-field region our atoms will seek. Symmetry also ensures that the gradient along the axis of the coils, z, is twice as large as that in the x, y-plane. Near that zero, the field increases linearly with distance [170] and is written

$$\boldsymbol{B}(x, y, z) \simeq b' \left(-x, -y, 2z\right) , \tag{2.6}$$

where b' is the magnitude of the gradient of the magnetic field in the x, y-plane and is proportional to the current. The coordinates x, y, z are measured relative to the trap center at (x_0, y_0, z_0). Combining this with equation (2.5), which depends on the magnitude of equation (2.6), we can write the potential felt by an atom at $\boldsymbol{r}$ near the zero of a quadrupole field

19. This separation works out to be approximately 1.25 times the radius of the coils [170].

as

$$V(\boldsymbol{r}_a) \simeq \kappa\sqrt{x^2+y^2+4z^2}\,, \tag{2.7}$$

where $\kappa = g_F\mu_B m_F b'$. This will form a linear trap for atoms in low-field seeking states near the magnetic zero. Very near this zero the trapping force becomes weak, but of course the atom cannot leave the trap without passing through a region where it becomes stronger. However, the spin can undergo non-adiabadic transitions to other, untrapped states. These are the *Majorana losses*.

Majorana losses

The potential described in equation (2.5) relies on the on the atom moving in the magnetic field adiabatically, that is to say, changes in the magnetic field the atom experiences are slow compared to the internal dynamics of the atom. If this is not true, m_F and therefore $\boldsymbol{\mu}$ cannot me considered constants of the motion and the Hamiltonian (2.3) becomes time-dependent and the steps leading to a trapping potential (2.7) no longer hold. If an atom moves through a region where the magnetic field is zero, this is more-or-less guaranteed to happen since the quantization axis goes away and m_F becomes ill-defined. When it renters a region where there is a field, m_F is not conserved and the atom may be in another spin state. E. Majorana described the probability of these kinds of transitions when an atom moves *near* a region of magnetic zero in 1932 [171], hence the name. If the atom experiences a magnetic field weak enough, and changing quickly enough it can undergo one of these transitions to another spin state which constitutes a *spin-flip*.

Because we are looking for a breakdown in the adiabadic condition, we are looking to compare two time scales: one for the internal response of the atom to the magnetic field, given by the Larmor period, and one for how it experiences changes in the magnetic field. The Larmor period is just the inverse of the Larmor frequency [172], $\nu_L = \mu B(\boldsymbol{r})/h$. In our quadrupole trap, the field strength, $B(\boldsymbol{r})$, increases linearly with r', the magnitude of $\boldsymbol{r'} \equiv (x, y, 2z)$, like b'. This means, that the Larmor period of an atom with magnetic moment of magnitude μ a distance r' in the trap is

$$\tau_L(\boldsymbol{r'}) = \frac{h}{\mu r' b'}. \tag{2.8}$$

This is the period it takes the magnetic moment of the atom preceses around the magnetic field line at $\boldsymbol{r'}$ and serves as a time scale for the internal dynamics. If this is fast compared to the time it takes the atom to see the direction of the magnetic field change, τ_B, the adiabadic condition is satisfied. This can be assumed to hold unless we are close to the trap center where the field becomes weak and therefore τ_L becomes slow.

At the point where a τ_L is slowest and τ_B fastest, and a flip is most likely, the atom will be going at a tangent to a surface of equipotential. So, in the frame of reference of the atom, the only component of the magnetic field which changes is the one along its direction of travel [171]. This makes the rate at which it sees the magnetic field change orientation linear with the atom's velocity but inversely proportional to the distance from the center since the magnitude of all the components of $\boldsymbol{B}$ increase linearly with distance from the magnetic zero. This gives $\tau_B(\boldsymbol{r'}) = r'/v$, where v is the speed of the atom in the frame of the trap center.

The ratio of these two periods will give the adiabaticity condition. This reads

$$\Upsilon_{\mathrm{Maj}}(\boldsymbol{r}, t) = \frac{\tau_B(\boldsymbol{r'})}{\tau_L(\boldsymbol{r})} = \frac{hv}{\mu b' r'^2}\,. \tag{2.9}$$

When this reaches unity, Majorana spin flips become very likely.

To calculate the exact probability of any particular spin transition, the overlap of those quantum states must be calculated. For traps containing a mixture of spin states, a spin-change need not imply a loss and a more complicated expression than (2.9) must be calculated [173]. Since we have only one state in the trap, and the transitions allowed from it are to untrapped states, (2.9) is sufficient for estimating of a loss due to a non-adiabatic spin change.

2.2.5 Current imaging setup

Currently, our imaging setup is designed to take absorption and florescence images of the cloud in or released from the MOT or magnetic trap. Looking through the pictured porthole of the MOT chamber (marked (d) in Figure 2.3), one is taking the point of view of our camera and looking into the imaging beam. Details about the camera and lens are given in Chapter 4. As the science chamber comes online, imaging systems will be installed there including absorption imaging setups and, eventually, a light-sheet imaging system similar to the one described in [174]. In anticipation of that, I removed the higher-precision PCO "Pixelfly" CCD camera[20] [1, § 1.3.3] from the setup and integrated an ImagingSource CMOS camera.[21] The Pixelfly has pixels 3.75 µm to a side, a true CCD charge-counting feature and a double shutter feature that allows two images to be taken 1 µs apart. It will be used with the lenses described in Chapter 4 to do absorption imaging in the science chamber while the ImagingSource camera has pixels 4.4 µm to a side, no such extra features and will be used to image clouds in the MOT chamber.

I integrated the ImagingSource camera with our current MATLAB-based experimental control software using the MATLAB Image Acquisition toolbox. This toolbox provides controls for the camera parameters such as exposure, frame rate, signal gain and color correction. It is very important to suppress any kind of color correction since it can lead to the signals from some pixels being artificially enhanced, even when the camera is monochrome since the Image Acquisition toolbox will treat the signal as if it came from a color camera, even when the camera is monochrome and the software is told to use a grayscale color format. This means any color correction performed by the software will distort the signal.

We are trying to count photons and use that to deduce the number and distribution of atoms in the image (in a manner I will explain shortly) so *gain* added by the software should also be kept to a minimum. Gain is strictly defined as the ratio between the number of electrons collected by the sensor and the number it outputs as a "count" [175, §2.2.7]. What is given to the user to set under the label *gain* or *gamma* is a number, ϱ (I have used γ elsewhere), which relates the the signal, S like [175, Eqn. 2.21]

$$S = A\left(\int_{t=0}^{t_{\mathrm{Exp}}} I(t')\,\mathrm{d}t'\right)^{\varrho}, \tag{2.10}$$

where t_{Exp} is the exposure time, I is the intensity falling on the sensor and A is proportionality constant. If ϱ is far from unity, more work is necessary to obtain an atom number per pixel from S. There are methods [176], but they rely on knowing the ratio of the number of electrons excited to incident photons which is not readily accessible on cameras like the one from ImagingSource which do not provide a photoelectron count, although methods exist to extract it [92, 176]. The utility of minimizing the post-processing by the camera control software holds regardless of what kind of light signal is collected, fluorescence or absorption.

20. Model pco.pixelfly usb.
21. Model DMK 42BUC03.

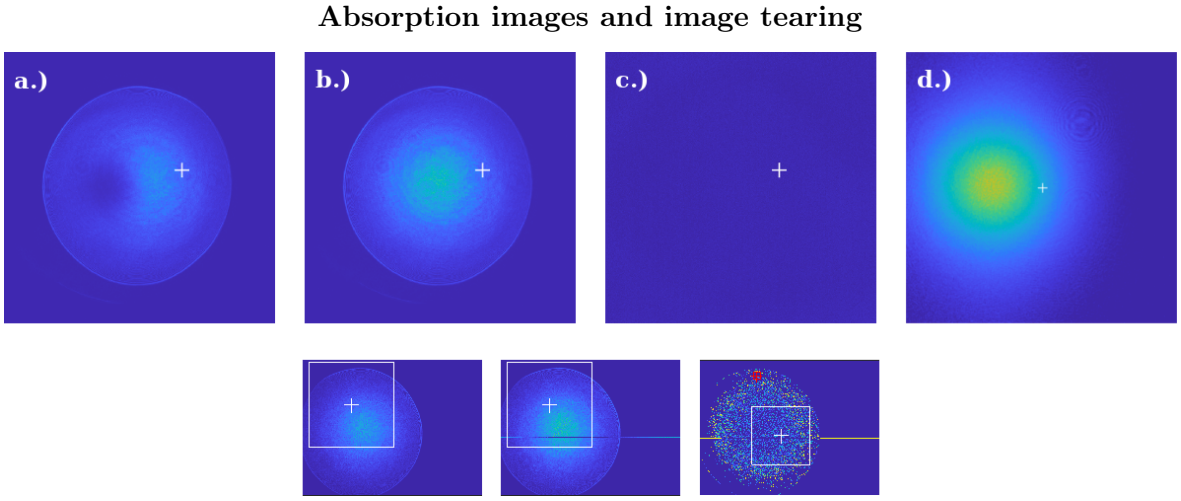

Figure 2.5 — Top row — Typical absorption image of a cloud of 1.88×10^8 atoms 6 ms after being released from the magnetic trap. Individually these are a.) absorption image of the atoms, $I(x, y)$ in (2.14) b.) an image of the laser light without the atoms, $I_0(x, y)$, c.) a background image for correction, $I_{\text{dark}}(x, z)$, and d.) the final, corrected and cropped image representing $n(x, y)$. Bottom row — example of the imaging fault described in § 2.2.5.2 appearing in an image of very few atoms. From left to right: with atoms, no atoms, and the corrected image. Notice how the a line of pixels in the background image seems shifted.

Fluorescence imaging works by putting the atoms in the excited state and then collecting the light they emit when they decay. This means we are only collecting the fraction of the light emitted into the solid angle subtended by the lens so the signal tends to be weak unless there are a large number of atoms and plenty of light, as in the MOT, when we use it for verification and centering the magnetic trap bias. Our data about the atomic cloud comes from absorption imaging. For details about fluorescence imaging, apart from the light-sheet imaging method, see [1, §1.4.1].

2.2.5.1 Absorption imaging

Absorption imaging obtains the atom number and integrated distribution from an image of the shadow of the atoms by taking advantage of the cloud's constant absorption cross-section in the low-saturation regime, where $I \ll I_{\text{Sat}}$. In this regime, the atoms can be relied on to be in the ground state because there not enough photons around to significantly populate the excited state. Therefore the likelihood of an atom interacting with any photon is just the absorption cross-section of the transition, σ. This means that the likelihood of a photon being absorbed in any particular part of its path, $\mathrm{d}y$, is equal to this cross-section times the probability that it meets an atom in $\mathrm{d}y$ which is just $\mathrm{d}y$ times the linear density of the atoms. This is exactly the case described by the Beer-Lambert law. Extending this conception from a single photon to a beam of intensity I means the linear atomic density is replaced by the local atomic density, $n(x, y, z)$ and we must take into account the intensity itself, since the more photons there are, the more likely they are to be lost. This makes the change in the intensity over the path length $\mathrm{d}y$

$$\mathrm{d}I\,(x, y, z) = -\sigma I(x, y, z) n(x, y, z)\,\mathrm{d}y \tag{2.11}$$

where the cross-section is given by

$$\sigma = \mathcal{C}\frac{\sigma_0}{1+\frac{4\delta^2}{\Gamma}^2} \tag{2.12}$$

where $\mathcal{C}$ is the sum of the squares of the Clebsh-Gordon coefficients of the accessible transitions, δ is the detuning of the laser from the transition with resonant cross section σ_0 and natural line width Γ. All of these are either set by us or obtained from [135]. Equation (2.11) suggests that I is an exponential function. Integrating over y would give the intensity distribution as a function of the initial intensity distribution, $I_0(x,z)$, and by taking the natural logarithm we can get an expression relating the initial and final intensities to the integrated density profile. This reads

$$\ln\left[\frac{I(x,z)}{I_0(x,z)}\right] = -\sigma\int_{-\infty}^{\infty} n(x,y,z)\,\mathrm{d}y\,. \tag{2.13}$$

The right-hand side is some function in x and z which we can obtain by measuring the intensity distributions on the left-hand side.

In practice we do not just measure $I(x,z)$ after interacting with the atoms and $I_0(x,z)$ of our imaging beam not interacting with any atoms. There are all sorts of noise sources that can affect the image such as the ambient light or even noise on the CCD itself [175, §7.1]. In order to counter act this we take an additional image, in the absence of atoms *and* light giving an "intensity distribution" of this noise. In fact, thanks to the Pixelfly's double shutter function we take this image twice, 1 µs apart. With the ImagingSource installed we still take two such pictures, since the treatment software expects them, albeit they are taken further apart in time. Call the "intensity distributions" of the background noise $I_{\mathrm{dark}}(x,z)$ and $I'_{\mathrm{dark}}(x,z)$. Then we calculate

$$\frac{1}{\sigma}\ln\left[\frac{I_0(x,z)-I_{\mathrm{dark}}(x,z)}{I(x,z)-I'_{\mathrm{dark}}(x,z)}\right] = \int_{-\infty}^{\infty} n(x,y,z)\,\mathrm{d}z \tag{2.14}$$

to arrive at the two-dimensional density profile. Integrating this over again gives a on-dimensional profile and a final integration gives the total atom number, N. In the experiment, these integrations about to performing discrete sums along all of the rows or columns.

In order for the distribution to have meaning, the intensity profiles cannot be distorted by gain being non-linear with respect to the intensity of the light or by color correction changing the response of the pixels unevenly. Thus we eliminate all of these corrections, leave the shutter open to receive the laser and eliminate any problems with over-exposure by reducing the power in our imaging beam.

2.2.5.2 Imaging faults

When I had integrated the ImagingSource camera, there started occurring experiments where our atom number would drop to near zero in between instances of it being $\gtrsim 10^7$ with little or no change in parameters. One of the images in those case would often appear blank, as if we had a triggering fault. Looking closer, some of the faulty images looked like those in the lower register of Figure 2.5, with rows of pixels shifted or reading zero. This is at least superficially similar to *screen tearing* where parts of one frame are drawn on top of another and it is usually associated with overtaxed memory resources [177, §3.3]. The standard Microsoft "Resource Monitor" tool on the computer controlling the camera showed large spikes in memory usage

by our manufacturer-supplied laser control software coinciding with faulty images. Since this software is only needed to change the laser parameters we can run the experiment without it open and these image faults do not occur.

2.2.6 Transport chain

From the magnetic trap in the MOT chamber the atoms can be loaded into the transport chain. The optimization of this is the focus of Chapter 5 so interested readers are referred there but I will mention a few technical details about the control mechanisms here.

The current to magnetic transport is supplied by three Delta Elektronika power supplies,[22] capable of supplying zero to 15 V and up to 110 A. These are governed by electronic switching boxes, built by the electronic workshop at LPL. Exact schematics for these switch boxes can be found in [1, App. D]. The cables from the magnetic transport, including the MOT chamber are connected to three switching boxes that direct the current from the power supply to the various pairs of cables that are attached to them. They receive a signal from a fourth, control box which tells them which cable pair to direct the current to. When the control box receives a TTL voltage-high pulse, it switches which configuration of switch boxes are supplying current to which of their cable pairs. In this way we can control which coils in the transport chain are receiving current and so move the magnetic trap (see § 5.1.2). Magnitude of the current is set at the power supply.

The concept of using a right-angle magnetic transport to isolate parts of a cold-atoms experiment is not original (we base our design on that of Greiner et al. [2]) but it is crucial for us to achieve a high repetition rate and make studying fluctuations practical. It also ensures that the MOT and the science chamber do not share any axes so they can be optically accessed independently of each other. Flexible bellows were necessary in both legs of the transport to compensate for small misalignments between the chambers and the corner. The valve marked "V" in Figure 2.2 is just after the corner (reckoning along the path of the atoms on their way to the science chamber) and kept the first leg of the transport isolated from the second leg and the science chamber. This is what allowed us to optimize the transport before the science chamber was installed and while the chip and its mount were being worked on.

2.2.7 Science chamber and chip enclosure

The chip and its mount, where we will perform the final evaporative cooling and trap the quasi condensate, are the focus of Chapter 3 so here I will just mention a few things regarding its assembly and the vacuum difficulties which allowed me to do the work in Part II.

The science chamber itself is an octagonal chamber manufactured by Astemec [1, §1.2.3]. On the bottom is a CF63 opening for a window through which we intend to do light-sheet imaging as in [178]. The chip is mounted through the top with another CF63 connection. After the chip had been mounted and tested as I describe in Chapter 3, it and its housing had to be mounted on top of the science chamber. We accomplished this with the help of a *jug sling hitch* [179, #1142].

There is a CF16 connection to the magnetic transport on one side and then seven CF40 connections, one for the valve and NEXTorr ion pump[23] assembly and six for silica windows. All the windows in the science chamber were originally sealed with HELICOFLEX® gaskets,[24]

22. Model SM15-100.
23. Model D 100
24. Model HNV200, a "Delta" model with small raised edges in the center of its faces where it make contact.

one on each side of the window, sealing to either the chamber or the flange. These gaskets are springs with metal sheathes around them that push back against the window and flange ans so provide a tight seal [180]. These types of seals have been used successfully [181] and unsuccessfully [182]. After baking the chamber a fairly low temperature ($< 130\,^{\circ}\mathrm{C}$ so as not to cause a glass transition in PEEK material in the chip tower), we could not get the pressure on the ion pump down even to $\sim 10^{-7}\,\mathrm{mbar}$ whereas we could maintain a pressure below $10^{-10}\,\mathrm{mbar}$ in both the science chamber and the chip housing prior to assembly and baking. This was the start of a leak-hunting effort that took nearly a year, which involved many days of changing the temperature at less than a few degrees per hour so as not to crack the windows. In the end, we found that the HELICOFLEX® gaskets on three of the CF40 windows were leaking and we replaced all the HELICOFLEX® gaskets with indium seals [183, p. 272]. I can offer here a few small clues found searching for these failures.

First is that while cooling the chamber, and plotting temperature and pressure on the same axis, it could be seen that during some periods where the temperature decreased the pressure increased. This means that it cannot be an out-gassing phenomenon of any kind and suggests it is a mechanical leak; the gap opens as the parts cool and contract. Second is that small gaps ($< 1\,\mathrm{mm}$) could be observed between the inside HELICOFLEX® gaskets and windows that, in the end, where determined to be leaking. The gaps where barely visible to the eye, but could be seen with an $f = 50\,\mathrm{mm}$ singlet lens. They were small enough that they would not have appeared significant if not compared with the seals on the windows next to them. Lastly is to make sure there are no leaks in the pumping system when using a helium leak detector. We have a Varian Sentorr turbo pumping group comprising a Pfeiffer[25] primary pump and a Varian turbo pump[26] capable of reaching pressures as low as $10^{-8}\,\mathrm{mbar}$. We decided to use our Adixen[27] helium detector as a primary pump. This would let us search for leaks at a lower pressure and hopefully find smaller leaks. The only leaks that gave convincing signals ($\gtrsim 10^{8}\,\mathrm{ppm}\cdot\mathrm{L}^{-1}\cdot\mathrm{s}^{-1}$) when He was released near them where the leaks in between the turbo pump and the primary pump. When He was released near the chamber the helium detector never measured much higher than $\sim 10^{7}\,\mathrm{ppm}\cdot\mathrm{L}^{-1}\cdot\mathrm{s}^{-1}$ and the rise never happened in significantly less than 30 s. We hypothesize that these leaks between the turbo pump and the leak detector let the He escape because once we hooked the leak detector directly up to the science chamber the leak detector gave readings of $\gtrsim 10^{9}\,\mathrm{ppm}\cdot\mathrm{L}^{-1}\cdot\mathrm{s}^{-1}$ within seconds when He was released near the faulty windows. Even more decisive readings were obtained by using He to inflate a latex glove taped around the window flange. With these new seals, the chamber is holding at better that $10^{-10}\,\mathrm{mbar}$, the lowest pressure measured by our ion gauge, for many months.

2.3 Measurements to collect

All of the preceding is to describe a machine that traps and cools atoms for confinement in an effectively one dimensional trap on an atom chip. In this section I will describe how we intend to collect the data that will give us access to the physics described in Chapter 1. The main thing we will use is density profiles, captured by absorption as well as light-sheet fluorescence imaging. This is what we will use to study fluctuations and calculate the second order correlation function (equation (1.41)) that can be used to characterize the different

25. Model XtraDry.
26. Model Turbo-V 70.
27. Model ASM 142D.

phase regimes (see § 1.2.4). These kinds of imaging techniques give us access to the number of atoms in the trap and therefore loss rates, which should increase when the scattering length diverges at and near the Feshbach resonance we intend to demonstrate. Then there is Bragg spectroscopy by which we will probe the excitation spectrum.

2.3.1 Density profiles

Density profiles are so crucial for measuring fluctuations because they not only let us calculate the $g^{(2)}$ function, if they are taken after a time-of-flight, they give us access to phase information about the trapped gas. This way we can compare both phase and density fluctuations to identify, for example, the quasi-condensate phase where density fluctuations are suppressed bringing $g^{(2)}(\boldsymbol{r}, \boldsymbol{r})$ close to unity but phase fluctuations allow phase coherence over a length ℓ_ϕ (1.46). We will use three kinds of density profiles: time-of-flight profiles to access phase information about the trapped gas; *in-situ* profiles of the trapped gas; and light sheet fluorescence imaging which can give us single-atom resolution of the cloud after the time of flight [174].

2.3.1.1 Time-of-flight imaging

A way to gather information about the trapped gas' momentum distribution is to suddenly switch off the trapping potential and allow the gas to expand. Because the gas is in the transverse ground state of the trap, it will have a broad momentum distribution in this direction due to the Heisenberg uncertainty principle. That means it will quickly become much larger than the trapped gas and the initial spatial distribution will be dwarfed by the distribution acquired by expansion and the initial spatial distribution will be lost in the image. In the longitudinal direction, the velocity distribution is related to the phase in different parts of the gas. This lets us turn phase fluctuations into density fluctuations which can be captured with absorption imaging [88].

When the excitations in the trapped gas are of low energy (i.e. less than μ), they are phonons and the local velocity of the gas goes as the gradient of the phase like [73, Eqn. 6.16]

$$v(z) = \frac{\hbar}{m} \boldsymbol{\nabla} \phi(z). \tag{2.15}$$

When the trap is switched off, this becomes the initial velocity distribution. This will be modified by interactions, but in the z direction as soon as the particles are far enough apart, the influence of interactions falls off within a timescale on the order of $1/\omega_\perp$. This destroys information about fluctuations up a scale of $\xi\mu/\hbar\omega_\perp \sim 100\,\mathrm{nm}$, a scale we do not hope to resolve. This means that information in the z direction is largely preserved. If the gas were not strongly confined, interactions could appreciably affect the velocity [184]. From that time on, the gas freely expands. This means that phase fluctuations translate into different local velocities which then translate into density fluctuations, appearing as stripes in the final absorption image [88, 184].

This means that these time-of-flight images will contain information about the correlation length, ℓ_ϕ [185]. The time-dependent correlation function reaches an asymptotic state with fluctuations on the scale of ℓ_ϕ in a time $t \sim m\ell_\phi^2/\hbar$. This is usually a short amount of time, 10 ms corresponding to a coherence length of 5 µm (Cf. [184]). In the quasi-condensate regime, these fluctuations are expected to be significant (see § 1.2.4).

2.3.1.2 In-situ imaging

The most straightforward way, at least the way requiring the least deduction, to get the density distribution of the trapped gas is to take a picture of it [92]. This is what the lens conversion discussed in § 4.3.1.1 was designed for. This is a method of measuring the density fluctuations used to calculate the $g^{(2)}$ function we will use to distinguish different regions of the phase diagram.

At first glace, it is a straightforward absorption image taken of the long side of the gas. In practice, it presents a number of experimental challenges over time-of-flight imaging. The method of obtaining atomic distribution (2.14) assumed that the cloud was not optically dense and that all the atoms along the path of the light have a chance to interact with it. In the trap, this assumption may fail and then one needs a method to calculate an *effective* absorption cross-section. Obtaining the effective cross-section involves probing the cloud with beams of different intensities, below and above the saturation intensity ($9.39\,\mathrm{mW}\cdot\mathrm{cm}^{-2}$ for the sodium D_2 line [135]). The maximum optical depth of the sample is then calculated assuming a different effective cross-section. In an ideal world, the correct effective cross-section would give a maximum optical depth that did not vary with intensity. In the real world, a working effective cross-section can be found by minimizing the variance in the maximal optical depth as a function of intensity [186].

2.3.1.3 Light-sheet imaging

Both of these imaging techniques rely on forming an image on a CCD and so face the usual limitations on resolution that come from objectives and finite pixel sizes. A way to greatly increase the resolution is to use a light sheet, where, upon passing through, atoms will scatter photons which we will collect with a lens. This kind of fluorescence imaging has the potential for single-atom resolution [174]. As they will only scatter photons as they fall through the light sheet, and only a small part of them will be scattered into our lens, it is necessary to maximize the numerical aperture (NA) (see Figure 4.1) of any lens used for this kind of imaging. This has an advantage, however, since high NA lenses tend to have very shallow depths of field ((4.5) & (4.10)). If this shallow depth of field is matched with the thickness of the light sheet, any spurious signal from outside the light sheet will fall out of focus. In principle, wth a high frame-rate, we could use this to reconstruct the whole three-dimensional cloud [111]. But in a one-dimensional gas, the two transverse directions are the same, so reconstructing the vertical direction would not provide any additional information.

As a kind of fluorescence imaging, atoms emitting photons will recoil, undergoing random walks in momentum space. In order to prevent a drift in the atoms due to light forces, the sheet will be made up of two counter-propagating lasers of equal intensity, with perpendicular linear polarizations to avoid setting up a standing wave. A repumping side-band, as in the MOT will be added to keep atoms out of the dark state. To detect individual atoms, which will naturally be well below the resolution limit of our imaging device, the single-atom autocorrelation function is calculated. This discretized convolution of the corrected image signal, S, with itself reads[p. 12] [174]

$$A(\Delta x, \Delta y) = \sum_x \sum_y S(x,y)S(x+\Delta x, y+\Delta y) \tag{2.16}$$

where Δ denotes the inter-pixel spacing in that direction. While the image of a single atom will be spread over several pixels by our optics, as long as point spread function, which

determines the image of a point source created by an optical system, does not distort the image in such a way that the atom is not at the center (unlikely in a system with a round aperture), the atom will be at the location where copies of its image overlap. At that location, the autocorrelation function spikes (see [174, Fig. 5(b)]) and we have located the atom. This lets us obtain resolution beyond that which the optics are capable of.

The full momentum distribution, especially at single atom resolution, is especially interesting because it will have a high-momentum tail that will fall off like $\mathcal{C}/k^4$ where $\mathcal{C}$ will be Tan's contact [187]. This number contains much information about the interaction forces between the atoms [188,189]. While it was originally worked out in a box potential it has recently been characterized for harmonic potentials like those on our chip [190,191]. This is still a difficult measurement since the high-momentum particles are fewer and produce a weaker signal and "high"-momentum depends on the value of $\mathcal{C}$, which increases with the interaction strength. Since atoms are more likely to be lost at high interaction strength, this measurement poses some challenges.

2.3.2 Losses

Losses due to interactions are called *recombination* or *N-body* losses. This is where N bodies undergo an inelastic collision. Often, this results in a molecule and, typically, a free atom. The two lowest orders are the most common [32, § III.A.2]: two-body losses where a pair of atoms usually undergo spin-exchange as discussed in 2.2.4[28] and three-body losses where three atoms interact to form a dimer and free atom carrying away the excess energy as kinetic energy [193]. This kinetic energy is usually enough to expel the atom so all three atoms will be lost. There are higher-order recombination losses, but the rate of N-body losses falls off like $|a|^{3N-5}/(3N-5)(N-1)!$ [194] where a is the scatting length, typically in meters and much less than one, even when "large." These are all more likely with longer scattering length. They increase drastically at Feshbach resonances and can be used to locate these resonances [39,127,195,196].

Because we are using atoms trapped in a particular spin state ($|F=1,m_F=-1\rangle$ see § 2.2.4) and are inducing the resonance with photons of a particular polarization (see § 1.3.3) we can say which order is more likely. In the lowest possible state ($|F=1,m_F=1\rangle$ for sodium), two-body recombination will be forbidden and three-body losses will dominate [127, 195].

2.3.3 Bragg spectroscopy

In order to probe the excitation spectrum, particularly in the strongly-interacting regime which has not yet been directly accessed by experiment, we will use *Bragg spectroscopy*. It was originally developed as a method to probe the structure of crystals by scattering photons off the grating provided by the crystal structure [197]. In this context, it means scattering atoms off of a grating created by a standing wave of light [198].

The setup we will use is similar to that of [199] and consists of two beams of light with momenta $\boldsymbol{k}_1$ and $\boldsymbol{k}_2$, having an angle θ and frequency $\delta\nu$ between them. Since $\delta\nu \ll \nu_1, \nu_2$ we can take an average magnitude of the two momenta, $k = (k_1 + k_2)/2$, having associated frequency ν and use it for certain purposes.

28. Dimer formation is also possible in some species such as cesium which has more complicated concessional dynamics than lighter alkalis [192].

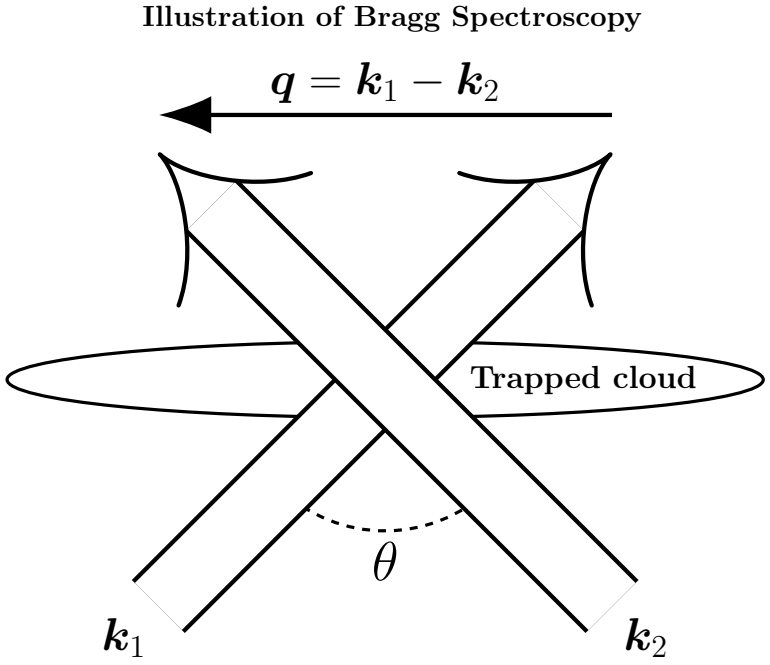

Figure 2.6 — Schematic of the two Bragg spectroscopy beams, having momenta $\boldsymbol{k}_1$ and $\boldsymbol{k}_2$ crossing over the trapped cloud at angle θ. An atom absorbing a photon from beam one and emitting into beam two will absorb a net excitation with momentum $\boldsymbol{q}$.

Atoms interact with the two beams by absorbing a photon from one of them and undergoing stimulated emission into the mode of the other. This makes an effective two-level system where the ground state is coupled to a particular excited momentum state (or rather *two* particular momentum states, which one depending on which beam it absorbs from) by a two-photon process. The momentum on the other hand, assuming the atom interacts with the beam having $\boldsymbol{k}_1$ first, will be $\boldsymbol{q} = \boldsymbol{k}_1 - \boldsymbol{k}_2$, or the other direction if they interact with the other beam first. Either way, by vector addition, the momentum added will have a magnitude of $q = 2k\sin(\theta/2)$ and a direction perpendicular to the bisector of the two beams since $\delta\nu$ is small. We can align this direction with the axis of our trap (see Figure 2.6). The energy imparted by this interaction will be $\delta E = h\delta\nu$. Because we are probing excitations close to the ground state, $\delta\nu$ should be small, justifying the condition mentioned before. If one wants to select a particular energy level to excite, the variability in $\delta\nu$ should be less than the gaps in the longitudinal excitation spectrum, $\hbar\omega_z$ in the harmonic case.

The condition that atoms must meet to be resonant with this two-photon transition is (the *Bragg condition*) that they have initial momentum k_i satisfying $h\nu = q^2/2m + \hbar q k_i/m$. This lets one select different initial velocity classes to address by adjusting θ and hence q [200,201]. It also lets one calculate the precision with which θ must be maintained such that variations in $h\nu$ are small compared to the separation between the excitations one wants to address.

We are interested in using it to probe the excitation spectrum of the trapped gas in different regimes. The excitation spectrum is defined by a *dispersion relation*, giving the energy of the excitation as a function of its momentum. The excitations induced via Bragg spectroscopy become fluctuations in the phase which become fluctuations in the density after time-of-flight as mentioned in § 2.3.1.1. This essentially knocks the excited portion of the gas out from the rest of the cloud and by measuring the atom number in this separated

portion, the total momentum and momentum per atom transferred from the Bragg beams can be determined [200]. By scanning the difference in frequency, $\delta\nu$, and hence the energy associated with this added momentum, the dispersion relation can be determined [202].

The excitation spectrum has been probed experimentally in three dimensions [199, 203], and in one dimension by Meinert et al. [204], however the one-dimensional experiment relied on an array of tubes. Mienert et al. used Bragg spectroscopy to excite excitations in an array of tubes which they then averaged over. The distribution of excitation energies was so wide, that it, as well as other results, were reckoned as evidence of the two branches of the excitation spectrum. We hope to excite these branches directly and independently and reconstruct the precise Lieb excitation spectrum [49]. This will not be trivial, since the excitations tend to be a mixture of the two branches of the Lieb spectrum, with a bias toward the Lieb I branch [205]. A perfect comparison will have to wait until we can make a box potential, but in our harmonic trap we will have to take into account the variations in density over the length of the trap [199].

2.4 Conclusion

The first goal of this experiment is to explore the phase diagram of a one-dimensional gas of bosons using a novel way of inducing a Feshbach resonance. The different regimes are characterized by differences in the correlations between their fluctuations from mean density. This already starts to dictate our experimental choices. We need a reliable way to cool and trap atoms, with a high repetition rate and a way to shine microwaves onto the trapped gas to create the Feshbach resonance. A Zeeman slower and a MOT are well-tested technologies and an atom chip can create the high gradients necessary for exploring one-dimensional physics and can incorporate a coplanar microwave guide for the microwave induced Feshbach resonance. A permanent magnet Zeeman slower even eliminates the need to supply coolant and electrical power. Separating the MOT chamber from the chip chamber increases the repetition rate, since we can load our MOT while we are conducting experiments on the chip while protecting the atoms on the chip from light scattered by the MOT. This also give us more optical access to the science chamber; all of the beams for the initial cooling and trapping will not be in the way. This is especially helpful for when we will set up our Bragg spectroscopy which requires a significant and adjustable angle in between the Bragg beams and an absorption imaging system on top of that. But this separation raises the question of how to get the atoms from the MOT chamber to the chip efficiently, hence the optimization of the magnetic transport discussed in Chapter 5.

The current state of the experiment is that the vacuum chamber is complete, including the chip chamber which is now holding a vacuum. We can load our MOT with $\sim 10^8$ atoms in $< 10\,\mathrm{s}$ and transport one third of those to science chamber in around $1.2\,\mathrm{s}$ without much heating (see Chapter 5). The chip is installed and next will be to load it, perform evaporative cooling, and look for the Feshbach resonance by observing atom loss. From there, we will explore the phase diagram, in particular the strongly interacting regime. We plan to probe the Lieb–Liniger excitation spectrum via Bragg spectroscopy. To make the comparisons to the Lieb–Liniger solution even more exact, we hope one day to construct a box potential so we can do away with approximations like the local density approximation to account for our harmonic longitudinal confinement. The idea would be to generate the strong, cylindrical confinement with the chip wires and then to use repulsive, blue-detuned lasers to form the ends, similar to how Meyrath et al. [206] formed the ends of their all-optical box.

Chapter 3

Atom chip with coplanar waveguide for a one-dimensional trap

In order to get a single, one-dimensional trap to compare with theoretical results we have elected to do our final trapping on an atom chip [45]. These use micro-fabrication techniques to form metal wires on a substrate. Passing current through these wires will generate magnetic fields that, when combined with external bias fields, will make steep gradients that can form the kind of tight confinement we want to reach low dimension. This kind of magnetic trap is compatible with a microwave Feshbach resonance where it would not be with a static magnetic field one because a large magnetic field could move the trap relative too the chip. If it moves it too far, the gradients become too shallow to maintain tight confinement. If it moves too close, roughness in the trapping wires can lead to irregularities in the trapping potential [207]. It can also be combined with optical trapping techniques to confine the atoms in the longitudinal direction and create a homogeneous trapping potential that will allow for more direct comparisons to the Lieb–Liniger model.

The chip itself is made of gold deposited on a silicon wafer via optical lithography techniques [45, Ch. 3]. It incorporates two trapping wires, one 100 µm and another of 50 µm wide, wires to deliver radio frequency (rf) oscillating fields to the trap, and a coplanar microwave guide, as well as some other wires necessary for trapping. This collection of wires should give us substantial control over the trapping potential and give the experiment a long working lifetime. The immediate goal is to look for the microwave-induced Feshbach resonance and explore the one-dimensional phase diagram. But the two trapping wires and radio-frequency lines give us more control over the trap geometry (e.g. [146]) offering us the possibility to explore other kinds of physics on the same apparatus.

The chip is fixed on a mount which must hold the chip at the correct position, carry current to the wires on and around the chip, and act as a heat-sink. Since the chip is under vacuum, all of the energy created by Joule heating must be carried away through the mount, into the vacuum chamber walls where it can be dissipated into the air. At the same time, shorts must be prevented; even the on-chip micro-wires are expected to carry in excess of an ampere at times. Vacuum compatibility must be considered at every step when designing and constructing the chip assembly. Vacuum compatible materials that do not out-gas at low pressure must be used and care must be taken that no pockets of gas get trapped between surfaces that can escape later when the pressure decreases.

Generating a magnetic tube trap with a wire

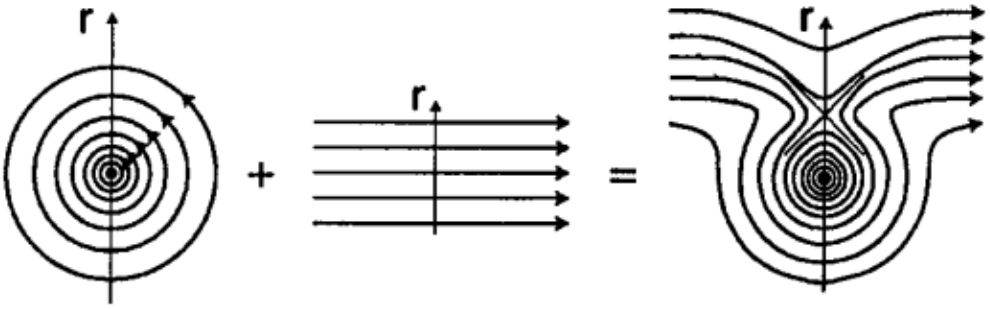

Figure 3.1 — Current in the wire generates a magnetic field around it, then adding the bias field creates a magnetic zero where low-field-seeking atoms will be trapped. This creates a tube with open ends that will have high non-adiabadic spin-flip (Majorana) losses (see page 105). Taken from [147].

The central content of this chapter is the results of the electrical tests performed on the chip assembly by the ingénieur de recherche of our group at LPL, T. Badr, who was instrumental in the design and fabrication of the chip and its mount, and myself. Therefore, in Section 3.1, I will go over the theory needed to evaluate the data. Then in Section 3.2 I will discuss the design and construction of the assembly in detail. Then in 3.3 I give our results. Finally in Section 3.4 I make comparisons between the data and the expected field strengths and draw my conclusions.

3.1 Theoretical background

Here I will go over the ideas needed to evaluate the data I gathered testing the electrical properties of the atom chip. This means describing how to determine the magnetic field strengths we are going to need. First I will go over how the trap is formed out of magnetic fields generated partially by the wires on the chip. There are also bias fields generated by wire coils outside the science chamber and the "Z"-wire under the chip that is used for loading the trap on the chip. For this section I relied on a 2002 review by R. Folman et al. [208]. Then I will talk briefly about the mechanism of evaporative cooling, since the chip mount incorporates the antenna we will use for this. Lastly I will talk specifically about the microwave guide that will generate the microwave field that will induce the resonance and make a few remarks about how the modeling was done.

3.1.1 Trapping with wires

As discussed in § 2.2.4, an atom with a non-zero magnetic moment will have a magnetic term like (2.3) and be subject to magnetic forces. As discussed in that sub-section, traps can be created with magnetic fields, which the chip will generate with wires.

A straight wire, carrying current I, at points far enough from the ends so that they can be neglected, will generate a magnetic field, $\boldsymbol{B}$, a perpendicular distance $\boldsymbol{r}$ from itself, having a field vector

$$\boldsymbol{B}(r) = \frac{\mu_0}{2\pi}\frac{I}{r}\check{\boldsymbol{I}} \times \check{\boldsymbol{r}} \tag{3.1}$$

where μ_0 is the permeability of the vacuum and $\check{\boldsymbol{I}}$ and $\check{\boldsymbol{r}}$ are unit vectors in the direction of the current and the distance from the wire respectively. This is illustrated in the leftmost part of Figure 3.1: equipotential rings going out from the wire, oriented according to the right-hand rule, with the field strength diminishing like $1/r$.

Low-field seeking atoms can be trapped by adding a bias field, $\boldsymbol{B}_b$ perpendicular to the wire, as seen in Figure 3.1. This will cancel out the field on one side of the wire, putting a magnetic zero at a distance $r_0 = \mu_0 I/2\pi B_b$ from the wire. This forms a two-dimensional quadrupole field with a one-dimensional zero and a field which increases linearly as one moves radially away from it. Low-field-seeking atoms will be trapped in this tube-shaped structure but can undergo non-adiabatic spin flips into untrapped states near the magnetic zero, as discussed in § 2.2.4. This can be avoided by a second bias field, $\boldsymbol{B}_{ip}$, along the axis of the trap which lifts the degeneracy of the trapped and untrapped states and so prevents non-adiabatic spin flips between them. This arrangement of fields is called an Ioffe–Pritchard trap [209,210].

This is the same kind of trap formed by the "Z"-wire (Figure 3.4, ①). This is a very heavy wire, the central part of which runs under and parallel to the trapping wires. Under the influence of the bias field it forms a two-dimensional quadrupole field, just as the trapping wire does, albeit farther away because the current through the "Z"-wire is much higher than the trapping wire. This forms an intermediate trap to take the atoms from the last three-dimensional quadrupole trap in the transport chain to the trap formed by the chip wires.

3.1.1.1 Forming the trap

There is still the problem that the wire is not infinitely long and the atoms can be lost out of the ends of this tube. Let me characterize the tube and its gradients before I address closing the ends to form a reliable trap.

In the tube, near the magnetic zero, the field will have a linear radial gradient [208, Eqn. 4],

$$\left.\frac{\mathrm{d}B}{\mathrm{d}r}\right|_{r_0} = \frac{2\pi}{\mu_0}\frac{B_b^2}{I} = \frac{B_b}{r_0} \tag{3.2}$$

where B_b is the magnitude of the bias field. Adding the field, B_{ip}, on the long axis of the trap, to eliminate the zero at the center make an Ioffe-Pritchard trap, shifts the potential so that is harmonic rather than linear. Now the field satisfies

$$\left.\frac{\mathrm{d}^2B}{\mathrm{d}r^2}\right|_{r_0} = \left(\frac{2\pi}{\mu_0}\right)^2\frac{B_b^4}{B_{ip}I^2} = \frac{B_b^2}{r_0^2 B_{ip}}. \tag{3.3}$$

Low-lying energy states, which are the ones we are interested in, can be approximated by a harmonic oscillator with ground-state frequency

$$\omega_\perp = \sqrt{\frac{\mu_B g_F m_F}{m}\frac{\mathrm{d}^2B}{\mathrm{d}r^2}} \propto \frac{B_b}{r_0}\sqrt{\frac{1}{mB_{ip}}}. \tag{3.4}$$

To make a trap with $\omega_\perp/2\pi = 2\,\mathrm{kHz}$, our intended confinement frequency, could be achieved, for example, with $B_b = 20\,\mathrm{G}$, $B_{ip} = 3\,\mathrm{G}$ and current in the trapping wire of $I = 1\,\mathrm{A}$. This would put the trap 100 µm from the wire (cf. [146] where it is 80 µm away). A lower current at the same bias field would tend to increase the trapping frequency since it will move the zero closer to the wire and increase the gradient.

There are several methods to close off the ends of the tube. For instance a set of coils can make a separate trapping field around the wire [208, §II.A.4.a]. The trapping wire itself can

be bent into a "U" or "Z" shape, for instance, with its ends going off in the same direction or opposite directions respectively [147]. In these configurations, the bent ends make a linear potential over the trapping region of the wire. In the "U" configuration , the field created this way reaches a zero between the wires where the fields made by counter-propagating currents cancel out. In the "Z" configuration, the fields add so there is no such zero. The bias field, B_0, tends to smooth out the potential made by the bent wire. It does this by moving down the elongated magnetic minimum which will move nearer the wire the further from the ends since the fields made by the end wires are circular (Cf. [147, Fig. 2]).

We use a slightly different method. There are four "U"-wires on the chip that come close to the trapping region, marked U1-4 in Figure 3.6. They have a portion on-axis with the trapping wires, whose field will not strongly affect the trap and a portion which is perpendicular to the trapping wire which will act like the bent portions of the trapping wires described above. Because they are separate from the trapping wire however, current can be supplied to them independently, and in either direction. This way we can make the kind of field usually associated with a "Z" configuration with our "U" wires, avoiding a zero in the field that closes the ends of the tube.

3.1.1.2 Loading the chip trap

Prior to be loaded into the trap on the chip, the atoms are held in the last quadrupole trap (see § 5.1.1) of the magnetic transport. This is a trap with a linear potential and a zero which should be at the geometric center of the science chamber which is $(3.37 \pm 0.10)\,\mathrm{mm}$ below the surface of the chip, or $(5.07 \pm 0.10)\,\mathrm{mm}$ away from the surface of the "Z"-wire. Once there the atoms need to be loaded into an intermediate, elongated trap formed by the "Z"-wire, a large wire made from oxygen-free high-conductivity (OFHC) copper with a rectangular cross-section mounted underneath the chip (labeled 1 in Figure 3.4). OFCH copper has less than 5ppm oxygen [211, alloy #C10100], is not volatile (so is suitable for ultra-high vacuum (UHV) environments), and posses excellent thermal ($\alpha \simeq 395\,\mathrm{W \cdot m^{-1} \cdot K^{-1}}$ [211, Tbl. 7.4]) and electrical ($\sigma \simeq 64.4\,\mathrm{MS \cdot m^{-1}}$ [211, p.8-4]) conductivity properties. This comes at the cost of being difficult to machine without deformation. So we contracted final machining to Aerosion, a company specializing in electrical discharge machining (EDM) [212], which removes material by vaporizing it by arcing electricity between the part and the tool, thus avoiding mechanical stresses which can deform the piece.

When current is run through the "Z"-wire and a bias field is applied, it creates an elongated Ioffe–Pritchard as described above. But the "Z"-wire is large, $(2.0 \pm 0.5)\,\mathrm{mm} \times (2.0 \pm 0.2)\,\mathrm{mm}$ at its narrowest point. This allows it to carry much higher current, allowing it to put the center of the trap much further away. A current of 80 A and a bias field of 32 G is enough to form a trap 5 mm away from the wire, just where the center of the trap formed by the transport chain should be. Then the transport trap can be shut off, and the current in the chip wires switched on. This will shift the center of the trap slightly toward the chip. The current in the "Z"-wire is then ramped down and the trap formed by the chip wires comes to dominate. If the ramp-down happens slowly enough, the atoms will follow the potential adiabatically and so be transferred from the large transport trap to the small chip trap.

This process can incur losses of hot atoms as the transfer occurs, especially if the maximum of the trapping potential is less than ten times the average temperature [208, IV.B.2]. The loss mechanism is hot atoms escaping the trap during the transfer so this will cool the cloud somewhat. The maximum energy the trap can hold is the *trap depth*, defined as the difference between the maximum and the minimum of the trapping potential. In this trap, the minimum

potential occurs at $B = 0$ and is zero itself. So the trap depth can be estimated by using the gradient, (3.2), the distance to the surface, and a factor of $|\mu_B g_f m_F| = \mu_B/2$ because $|g_F m_F| = 1/2$ in our case. At this energy, atoms will crash into the chip surface. This can be converted to a temperature by dividing by k_B. For the case of wires this will just be $\mu_B B_b/2k_B$ since the r_0 factors will cancel. But for the case of the "Z"-wire, where the surface of the chip is closer than the wire itself, this gives $\simeq 1.3\,\mathrm{mK}$ which is only 6.5 times a cloud temperature of 200 µK, so we can expect some looses this way. It is once the atoms are in the trap on the chip that we plan to cool them below this with evaporative cooling, but this process can begin in the quadrupole trap. It is not clear if the antenna we plan to use will be able to cool the atoms in the quadrupole trap as configured, however (see § 3.3.4).

3.1.2 Evaporative cooling

The technique of *evaporative cooling* was key to forming the first BEC [19]. It involves coupling high-energy states inside the trap to untrapped states [213]. We achieve this via resonant coupling with radio waves emitted by an antenna of copper wire coiled around the upper tower (3 in Figure 3.4). I will only address here how it is possible to cool atoms with radio waves and enough of the theory explain a few experimental choices. For fuller treatments see [214] or [75, Ch. 12].

The atoms are in a purely magnetic trap and so are subject to the Zeeman effect where the degeneracy between the different spin projections is lifted and they take on their own energy sub-levels. By addressing these transitions, we can flip the atoms from the trapped state, $|F = 1, m_F = -1\rangle$, to an untrapped one, $|F = 1, m_F \neq -1\rangle$, and eject them from the trap. By addressing larger gaps we can eject faster-moving atoms because these reach regions of higher magnetic field and so have larger splitting between the Zeeman sub-levels. In this way, the higher-energy atoms are coupled to untrapped states and the remaining atoms re-thermalize, they will do so a lower temperature. The photons that perform the coupling are typically in the radio frequency range with the exact frequency depending on the scaling of the Zeeman splitting with the field and the properties of the trapped cloud. If the frequency of the photons is ω_{rf} and the atoms are in a trap with a linear gradient, such as a quadrupole trap, they will make the trap depth $\hbar\omega_{rf}|m_F|$ [17].

Allowing the remaining atoms to re-thermalize is critical to the technique. For the reasoning that higher-energy atoms reach regions of higher magnetic field to hold, the gas must be sufficiently *ergodic*. This term will be defined in more generality in Chapter 6 but for now it means that the atoms' distribution in phase space depends only on their energy. This is a picture where the motion of the atoms is classical and the cloud re-thermalizes quickly as the atoms are ejected from the trap. These assumptions will be valid in the transport magnetic trap where the temperature is still much higher than the BEC transition temperature ($\approx 200\,\mu\mathrm{K}$ vs. $\sim 10\,\mathrm{pK}$ for a gas so dilute). So, in order to be sure that the rf-field is coupling out the atoms we want, we must start it at a frequency that couples out the hottest atoms in the trap and then ramp the frequency down over the course of tens of seconds, since the thermalization time in the initial, hot cloud is around three seconds (see 5.2.2) but will go down as the cloud cools and the phase space density increases.

Under weak magnetic field (a few hundred Gauss or less in the case of sodium in the $F = 1$ state) the energy Zeeman of the sub-levels scales linearly to first order [135, p. 7]. The Zeeman splittings between the $F = 1$ sub-levels is $0.70\,\mathrm{MHz} \cdot \mathrm{G}^{-1}$ [135, Fig. 2]. The transport trap has a horizontal gradient of $65\,\mathrm{G} \cdot \mathrm{cm}^{-1}$ and the size of the cloud is approximately 1 mm in this direction. This means the gaps we will be addressing are around 4.5 MHz and below

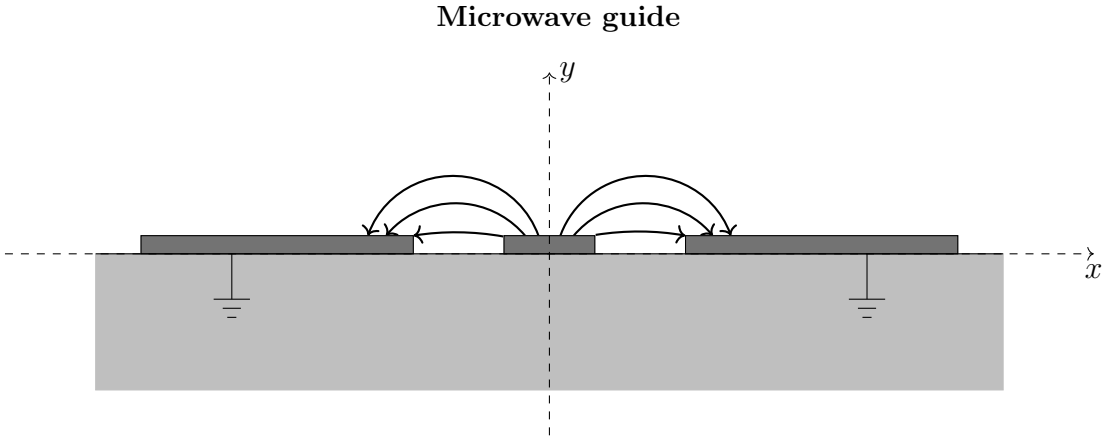

Figure 3.2 — End-on illustration of a microwave guide. Conductors in dark gray, insulating substrate in lighter gray. Arrows indicate the general shape and direction of the field lines. For a plot of the actual shape of the field lines see Figure 3.7.

(Cf. [17] where they swept from 30 MHz down to 1 MHz). We should be able to address at least eight times the gap or 36 MHz at around 100 mG.

3.1.3 Coplanar microwave guide

Once the atoms have been cooled to a point where τ (1.45) is at a point we are interested in, we can apply our other means of manipulation: the microwave-induced Feshbach resonance (see Section 1.3). The microwaves that will produce the resonance will be supplied by an arrangement of three of the wires on the chip that run straight down the center of the trapping region and form a *coplanar microwave guide* [46, Ch. 2]. A coplanar waveguide (CPW) is a type of *transmission line*, a structure designed to carry electricity, not just to get current from one place to another but also for the sake of the fields that current creates. A coplanar wave guide is a transmission line comprising a single, narrow strip of conductive material in between two wider conductors that serve as ground planes. These are set on top of a dielectric substrate, in our case the silicon wafer upon which our chip is constructed. When current is sent down the line, it creates fields extending to the ground planes on either side. A rough cross-sectional view is illustrated in Figure 3.2. If the current is oscillating, the field will oscillate, this is how we will get the microwaves for the Feshbach resonance onto the atoms.

A large part of the work in designing a microwave guide is determining its electrical characteristics such as impedance, capacitance, and coupling between the transmission line and the ground planes that might make for cross-talk. Current can be lost to the silicon substrat or through the substrate to the other wires [46, §8.4]. Skin effects, which are the result of the distribution of charge near the outside of the conductors in a shell or skin of finite depth, must be taken into account. Proximity effects, the effect the wires have on each other such as forming capacitors between them, must also be considered. Only after taking all of these into account can the fields generated by the currents be calculated [46, §8.2]. These calculations were performed R. Dubessy and resulted in Figure 3.7.

The method relied on the *Schwarz–Christoffel mapping* [215]. This is a *conformal* mapping, meaning it preserves angles, from polygons in the complex plane to another shape, including a straight line or open strip. Imagine taking the vertices of a polygon with some coordinated structure defined inside it and moving them to several points on a line; as in 3.3

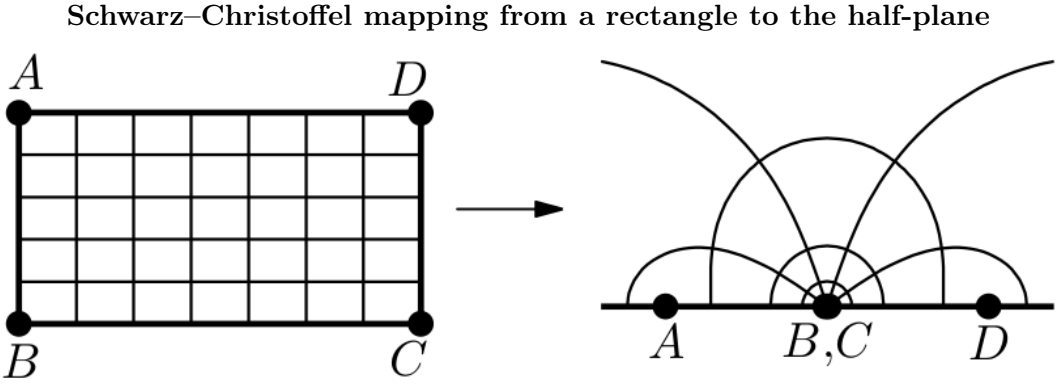

Figure 3.3 — An illustration of the Schwarz–Christoffel mapping that was used to estimate the field generated by the coplanar micowave guide (cf. 3.7). Here a rectangle, in which straight lines meet at right angles, is mapped to the half-plane with two corners are mapped to the same point. In order to map the field generated by the waveguide, the map would have to be to a piecewise constant line that matched the profile of the waveguide (see [215, §5.2]). Adapted from [215].

where the coordinate stucture is a simple grid. The Schwarz–Christoffel mapping will map the entire polygon to put the vertices at your chosen points and preserve angles between lines of equal coordinates. This is tremendously useful in fluid mechanical problems where lines of flux cannot cross and in electro-magnetic problems such as ours where lines of equipotential and field lines must meet at right angles.

To go from the rough illustration of field lines in Figure 3.2 to the precisely plotted shapes in Figure 3.7, one starts by taking the current as being carried in strips on the real axis of the complex plane. Then one applies the dielectric properties and boundary conditions of region (interior of the wires, vacuum above the substrate, dielectric substrate, or vacuum region containing the substrate) and gets the field lines. Then, one calculates the map to the rectangles corresponding to the cross-section of the wires. This takes the perpendicular filed lines and lines of equipotential from their grid-like structure in a homogeneous region (and the left side of Figure 3.3) to the curvilinear lines that meet at right angles as in the upper register of Figure 3.7 (or the right side of Figure 3.3). This mapping is necessary not only to determine the strength and geometry of the trapping fields, but is also useful for calculating the proximity effects on capacitance and resistance since it applies to the inside of the wires as well. For details see [46, §§8.2.1] and references therein.

3.2 Chip and support design and construction

Now that we have an idea of the kind of fields we want to make, let us discus the physical structures we will use to make them. There is the chip mount which includes the OFHC copper tower and rods which supply current to the "Z"-wire. There is the heat-sink the chip and sub-chip are glued to, the chip and sub-chip themselves and the various connectors and conductors that bring current between the chip and the outside of the vacuum chamber. All must be manufactured out of materials that will not give off gasses when placed under vacuum and careful attention must be paid to their thermal properties because the heat generated by the currents passing through the chip must be dissipated through the tower and into the stainless steel vacuum chamber housing. Electrical properties must be considered because there will be

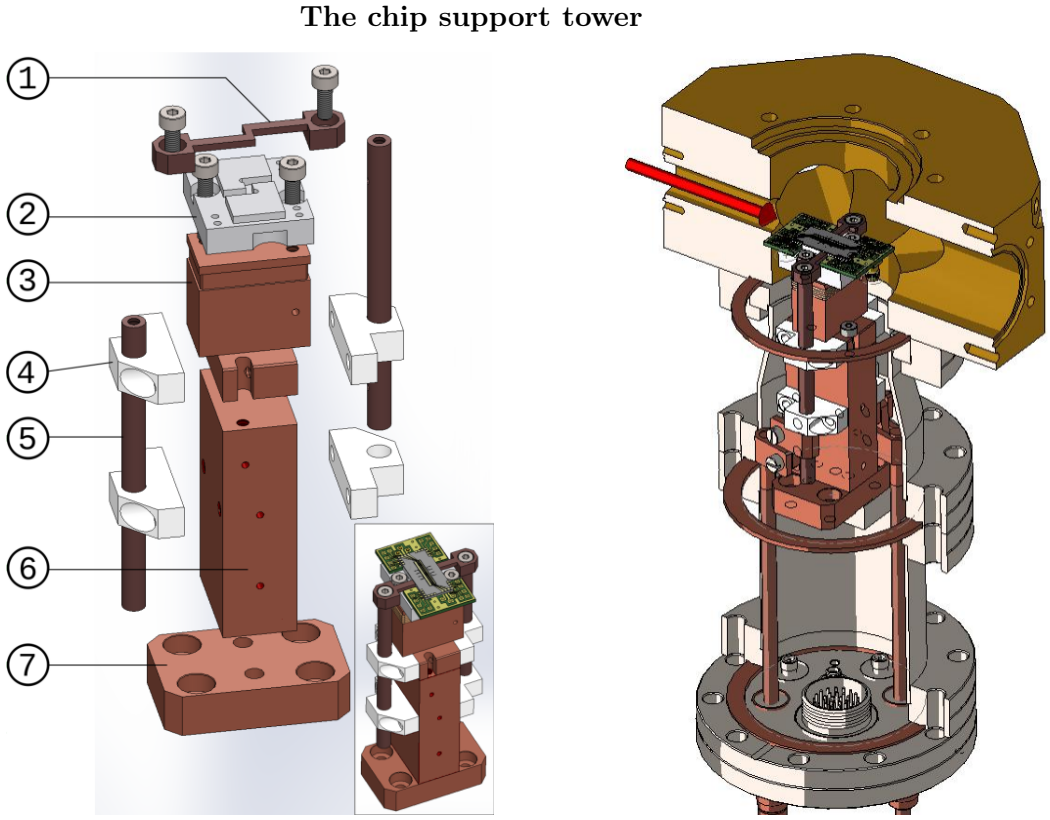

Figure 3.4 — Left — The chip support tower. ① Z-wire (OFHC copper). ② heat sink (Shapal^{TM} Hi M-soft). ③ Upper tower (OFHC copper). ④ Rod supports (PEEK). ⑤ Rod conductors (OFHC copper). ⑥ Lower tower (OFHC copper). ⑦ Base (OFHC copper). Right — Cutaway of the installed chip assembly. The red arrow indicated the path the atoms take to the chip. Wires are not shown but the full circumference of the copper gaskets in the CF joints is depicted. This tower is installed in the science chamber upside down. Images courtesy of T. Badr.

high currents which must be isolated from the large structures of high-conductivity copper. I will start with the housing and the chip mount comprising everything below the sub-chip including the heat-sink. Then I will describe the chip in detail and cover a few points about its construction and say few words about mounting it to the heat-sink and bonding it to the sub-chip.

3.2.1 Chip mount

The entire chip assembly is housed in a two-part stainless steel case held together by CF[1] joints with annealed copper gaskets. The electrical connections to the outside power supplies are made through a DN 75 CF end cap. The connections for the two copper rods (Figure 3.4, ⑤) which power the "Z"-wire (Figure 3.4, ①), pass through power pin feedthroughs from Kurt J. Lesker Co. welded into the end cap. These rods will carry 150 A. The evaporation coil and coplanar microwave guide are powered through SMA connectors, also from Lesker,

1. Conflat^{TM}.

welded into the end cap which feed lengths of high-frquency cable from Allectra,[2] rated up to 18 GHz. The cables and connectors are impedance-matched at 50 Ω to avoid reflections of signals at the connection [216, App. H]. The rest of the connectors for the other wires on the chip are fed through a 32-pin multisignal feedthrough from Lesker welded into the end cap. The wires that deliver current from this connector to the chip are single-strand Kapton^{TM}-coated 22 gauge copper. All wires are rated for 3 A.

The wires had to be connected to these feedthroughs before making connections with the chip, which meant that a certain amount of slack had to be left in the wires to make the final connection. To take up this slack without bunching the wires up around the chip tower, a 10 cm DN 75 CF stand-off pipe is used between the end cap with the connectors and the blind flange to which the base of the chip tower is bolted. Pieces of polyether ether ketone (PEEK), a polymer suited to UHV environments, are placed around the copper rods (Figure 3.4, ④) and the wires here to avoid electrical contact with each other and the cavities cut into the blind flange for them to pass through. This blind flange is where the heat generated by the chip and "Z"-wire is dissipated to the rest of the vacuum housing and so to the environment. It is from here that the distances must be exactly reckoned so as to place the chip just above the center of the science chamber where the last magnetic trap will be formed. How tall the tower needs to be depends on the exact length of the DN 75 CF to DN 63 CF reducer which measures (90.27 ± 0.10) mm from flange to flange.

To this blind flange is bolted the base of the chip tower (Figure 3.4, ⑦). This is broader than the rest of the tower to ensure good thermal contact. The tower itself is in three parts, a long, main part (Figure 3.4, ⑥), a spacer, and the top part around which the evaporation coil is wound(Figure 3.4, ③). The tower is constructed in parts because the machine which does the wire bonding can only have ≈ 20 cm between the working surface and the bonding tip. The chip and sub-chip must be glued to the ceramic heat-sink for the bonding to take place which means the "Z" wire must also be in place. The Shapal^{TM} ceramic the heat-sink is made out of is quite brittle and fragile so this should be mounted on a stronger material like the copper of the tower. If the evaporation coil is to be incorporated into the chip tower, it should be in the upper part so it can interact with the atoms. Thus the decision to construct the tower in parts, with a spacer that can be more easily machined/replaced to adjust the height. All parts of the tower are held together by screws. All screws, except for the four right by the chip are weakly-magnetic A316 stainless steel. The four by the chip are titanium to avoid stray magnetic fields. Care must be taken, especially with this kind of design incorporating multiple parts, that no pockets are left to trap air which could be released later when the assembly is put under vacuum.

The rods that deliver current to the "Z"-wire are 6.35 mm in diameter and are stabilized against by PEEK braces screwed to the side of the tower. They are connected to the rods coming from the feedthroughs by clips of OFHC copper ≈ 2 cm wide which are held tight by screws and can pass large amounts of current. Where the rods pass through the blind flange, they are shielded by PEEK sleeves. These PEEK components are one of the reasons we baked at such a low temperature. PEEK has a glass transition at 143 °C [217] and we did not want it to undergo this transition in the vacuum chamber.

The heat-sink under the chip is made from a ceramic, Shapal^{TM} Hi M-soft, which is electrically insulating but highly thermally conductive, having a thermal conductivity between 170 and 280 $W \cdot m^{-1} \cdot K^{-1}$ at 20 °C. The part has a channel for the "Z"-wire, screw holes and recesses for the SMA connectors on the underside of the sub-chip. The chip is glued to this

2. Type 380-SMA18G-MM-300.

The chip assembly

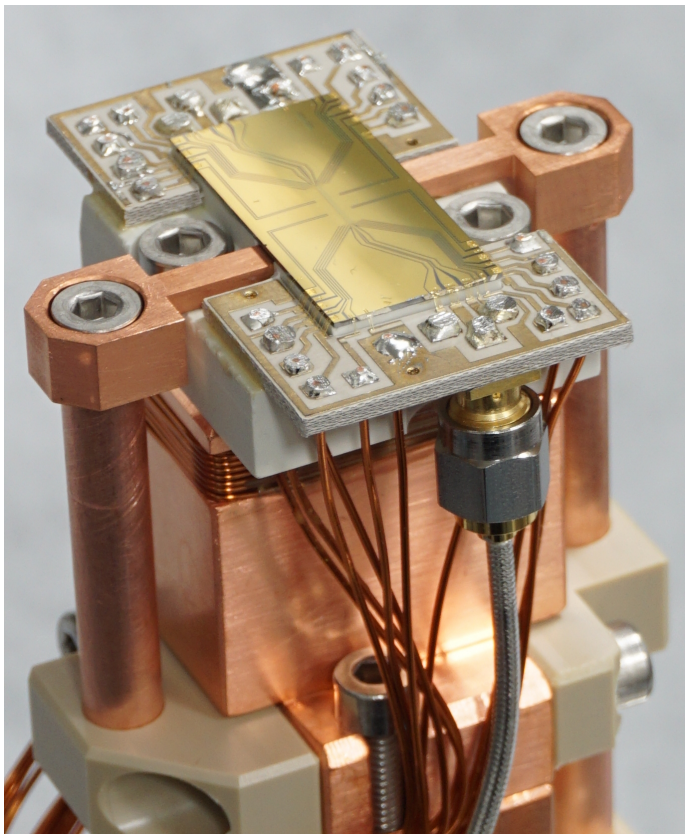

Figure 3.5 — Photograph of the elements discussed in this section. Compare with the chip and sub-chip schematic in Figure 3.6. Here the conductors soldered to the sub-chip and leading out of the vacuum chamber are clear.

with a two-component epoxy.[3] This has its own thermal conductivity of $0.89\,\mathrm{W}\cdot\mathrm{m}^{-1}\cdot\mathrm{K}^{-1}$. The sub-chip and the "Z"-wire are glued the heat-sink with a UHV compatible epoxy.[4]

3.2.2 Chip and sub-chip

Continuing the discussion of assembly from the last section I will deal with the construction of the chip and sub-chip and then detail the wires on the chip. This will lead into the electrical tests I conducted on the chip wires, the results of which for the content of the next section.

3.2.2.1 Construction

The chip itself (shown in Figure 3.5 and drawn in blue in Figure 3.6) is a $14 \times 29\,\mathrm{mm}$ piece cut out of a two-inch diameter, 500 µm thick silicon wafer, coated at the Technical University of Vienna by M. Stüwe. First, a 5 µm layer of photo-resist was applied by spin-coating. This resist hardens under exposure to ultra-violet light. Then a mask, which lets light pass where the wires should be is lain down and the sample exposed. The unexposed resist is then removed by chemical etching. Next a 20 nm layer of titanium is lain down by evaporation to

3. EPO-TEK 920-FL from FTPolymer.
4. 330-GLUE3 from Allectra.

improve bonding between the silicon and the gold. Then the 2 µm layer of gold is deposited. Lastly, the resist is pulled off, removing the gold from where it is not wanted. For a fuller overview of chip fabrication the interested reader is referred to [45, Ch. 3].

The surface quality and regularity of the wires must first be verified by optical microscope. This was done at the Université Paris Nord clean room, C(PN)2 with a Leica LWS3000 with a range of lenses from 5x to 100x and a numerical aperture of 0.9 and a resolving power of $< 1\,\mu m$. We had difficulty focusing on the edge of the tracks, indicating that they were trapezoidal rather than truly rectangular. This gold could have been removed when the resist was etched off as the resist fills gaps between the tracks and is etched off from the air side towards the substrate. That gold was removed somehow is also attested to by the fact that the spaces between the gold tracks were, on average, 2 µm wider than designed. Electrical tests (Section 3.3) confirm that we should be able to get enough current through them, but there is a reduced margin of safety.

Then the chip must be removed from the wafer. Silicon has a diamond cubic crystal structure [218, § 2]. This means that can be cut by *cleaving*. The wafer is scored with a scribe and then stressed such that it breaks along the scored line. The scoring must be done along a crystal plane, otherwise the wafer will shatter. In our case, the crystal orientation relative to the face of the wafer did not prove reliable so we resorted to sending the chips to GEM, an Austrian company partnered with the Technical University of Vienna, to be cut out with a micro-saw.

Since the wires running to the feedthroughs cannot be attached directly to the atom chip, we designed the sub-chip, depicted in green in Figure 3.6. It consists of two identical double-sided printed circuit boards (PCBs), designed by us and manufactured by PCB Electronics. Special care was taken in the fabrication choices to make it suitable for UHV applications. The substrate is not the typical one, which is often epoxy-impregnated paper or fiberglass, but AD350^{TM} from the Rogers Corporation, specifically designed for vacuum purposes. There is no resistive material laid over the wires or silkscreen printing done on them as these materials can out-gas. There is only a layer of nickel-gold laid on the 70 µm thick copper leads to prevent corrosion. The leads run between pads around *vias* (i.e. through-holes), where we can solder wires or SMA connectors, and 0.5 to 1 mm to a side pads by the chip where we can bond gold wires to carry current between the chip and the sub-chip. After the SMA sockets and wires are soldered in place, the ends are filed down so that nothing will interfere with the atoms on their way to the chip.

After the chip has been cut out, heat-sink, "Z"-wire and upper tower part can be assembled and the chip glued to the raised portion of the heat sink designed to accommodate it. Then the chip can be *bonded* to the sub-chip with gold wires. These wires have a $30 \times 13\,\mu m^2$ rectangular cross section and are attached to the surfaces via ultra-sonic welding. A tip containing the wire is manually lowered until it is brought into contact with the surface where it vibrates fast enough to melt the surface and the wire together. Great care must be taken not to press to hard or leave the tip in contact too long lest the surface melt away and no bond be formed. The tip is then manually raised and lowered on to the other surface where the other bond is made resulting in a connecting bridge 2 to 3 mm long. Since this is a manual skill, it requires much practice and we relied on the expertise of Thibaut Vacelet at LERMA, Observatoire de Paris to perform the final bonding.

Each micro-wire can carry up to 0.6 A, and since we want to be able to carry up to 2 A in some of the wires, multiple bonding wires must be used. As a rule, the wires on the chip, which have the strongest constrains on their size, should be the limiting factor when it comes

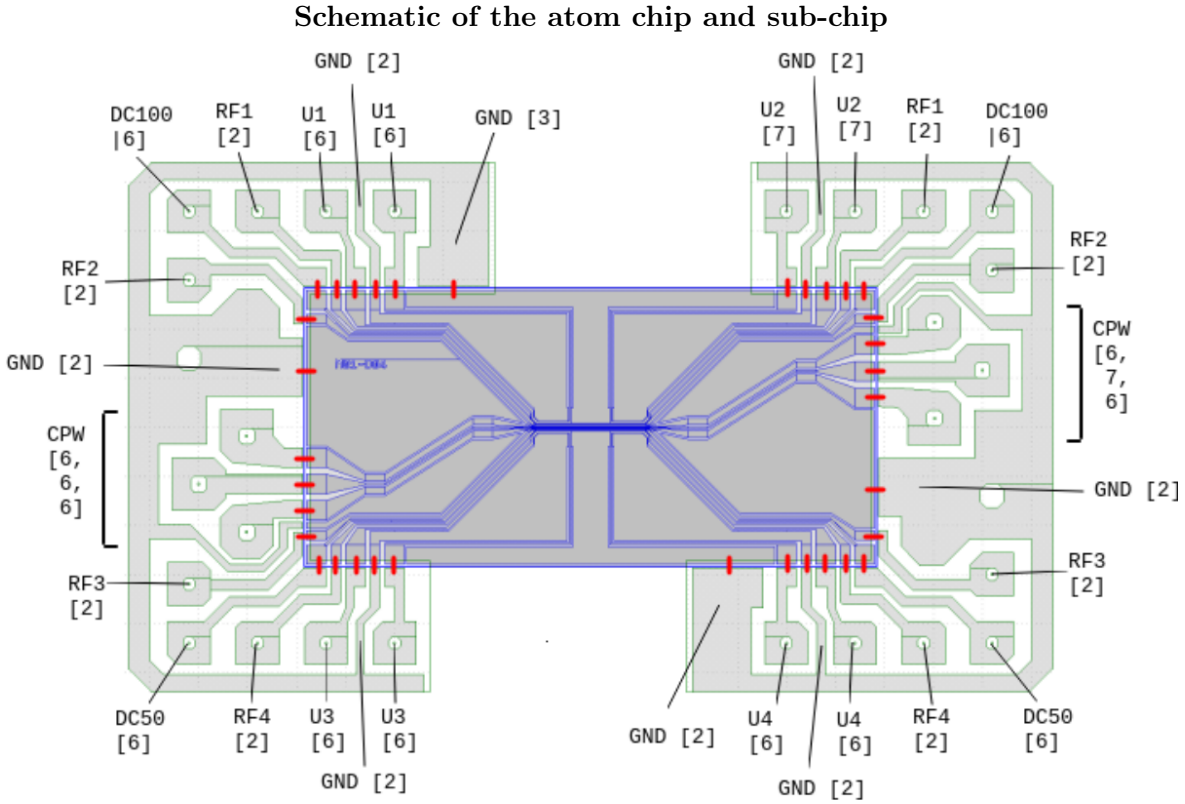

Figure 3.6 — Schematic of the atom chip (in blue) and the sub-chip (in green). Bonding between the two is indicated with a single red line. Labels on the sub-chip indicate the chip wire they connect to. Numbers in brackets are the number of bonding wires. Note how the "U"-wires pull away from the central trapping region where the wires run parallel in the order listed in § 3.2.2.2.

to current. This is why we use up to seven bonding wires for some of the connections, allowing up to 4.2 A to pass through the bonding wires.

All of these parts, as well as the vacuum chamber, were subjected to various kinds of chemical and ultra-sonic cleaning procedures that I do not think will be of interest to the reader. It suffices to say that no trace of any substance that will give off particles at low pressure should be left on any part that will be under vacuum. I will note that the micro-wires that bond the chip to the sub-chip are so fragile that even ultra-sonic cleaning may detach them. At that point one can only rely on gentile squirts of air or alcohol to remove dust or other impurities from the chip surface.

3.2.2.2 Chip wires

Here I give the name, abbreviation in Figure 3.6, designed width, measured width[5] if available, and describe the function of the different wires on the chip. Reckoning from the top of Figure 3.6 downward through the central trapping region, the wires are:

Radio frequency 1 (RF1) 10 µm (designed) — The two bonding wires can carry 1.2 A.

5. Measurements performed via optical microscope and software by Christophe Manquest at MPQ, Université Paris–Diderot

This is one of four radio frequency wires on the chip. These act as microstrip transmission lines [46, Fig. 3.5(d)] for delivering oscillating fields in the radio frequency range. These allow us to dress the trap or traps above the chip and create dressed trapping potentials [146, 219]. These wires can be thinner than the others since they are only expected to carry a few tens of milliamperes.

Very briefly, the idea is that the radio frequencies can induce spin flips in the atoms, changing the value of m_F and hence the trapping potential (2.5). If this happens quickly compared to the change the atoms see in the static magnetic field, the effective potential they see is a combination of those in different m_F sub-levels. See [220] for a review. This gives us an extra degree of control over our trapping potentials.

Trapping wire (wide) (DC100) 100 µm (designed) — The six bonding wires can carry 3.6 A

This is the larger of our two trapping wires that provide the direct current to form the traps as described in 3.1.1. Having two such wires not only gives us a back-up in the case of a failure, if the wires are too rough, the wider one will mitigate the effects of the roughness on the potential [221]. Furthermore, two trapping wires allow us to create different trap geometries, such as centering the trap over the CPW or creating double wells [145].

Radio frequency 2 (RF2) 10 µm (designed), 8.21 µm (measured) — Two bonding wires, 1.2 A.

Coplanar waveguide ground plane (CPW) 37 µm (designed), 34.76 µm (measured) — Six bonding wires, 3.6 A.

The ground planes of the CPW are there to pull the field lines to them as in the top register of Figure 3.7. As such they must be wide compared to the central microstrip so that any field lines not shaped by them will be beyond the region of interest.

Coplanar waveguide microstrip (CPW) 6 µm (designed), 3.58 µm (measured) — Limited by six bonding wires, 3.6 A.

This is the transmission line that will actually deliver the oscillating field to the atoms that will induce the Feshbach resonance.

Coplanar waveguide ground plane (CPW) 37 µm (designed), 35.18 µm (measured) — Six bonding wires, 3.6 A.

Radio frequecy 3 (RF3) 10 µm (designed), 7.79 µm (measured) — Two bonding wires, 1.2 A.

Trapping wire (narrow) (DC50) 50 µm (designed) — The six bonding wires can carry 3.6 A

Radio frequecy 4 (RF4) 10 µm (designed) — Two bonding wires, 1.2 A.

Then there are the "U"-wires that do not cross the trapping region. They run along the axis of the trap until they reach the region we expect the trap to be and then go at right angles to the central axis, thus taking the form they need to close the ends of the tube. The other bias fields will be generated with external coils.

"U"-wires 1–4 (U1–4) 100 µm (designed) — Six or seven bonding wires, 3.6 to 4.2 A.

The rest of the gold on the surface of the chip acts as a ground-plane, with two or three bonding wires making the connection. These are not expected to carry large amounts of current but it is an unnecessary risk to use just one micro-wire.

3.3 Electrical test results

Here I will summarize the results of the tests I performed on the chip wires and the "Z"-wire. More data, including oscilloscope readouts, can be found in Appendix B.

All of these tests were performed when the chip was installed in its housing and was under a vacuum of better than 10^{-10} mbar under the influence of our NEXTorr ion pump. The pressure never rose above this minimum readable value, indicating a lack of disastrous heating or out-gassing. This was necessary to test if the heat dissipation was sufficient. If the wires rise in temperature, T, their resistance, R, will rise according to the *temperature coefficient of resistance*, α_T, like

$$\frac{R(T)}{R(T_0)} = 1 + \alpha_T(T - T_0) \tag{3.5}$$

where T_0 is some reference temperature, such as room temperature, in Kelvin. We can detect this change in resistance by measuring the voltage applied to pass a certain current. For copper and $T_0 = 293\,\mathrm{K}$, $\alpha_T = 3.9 \times 10^{-3}\,\mathrm{K}^{-1}$ [222]. For gold it is $\alpha_T = 3.4 \times 10^{-3}\,\mathrm{K}^{-1}$ [223, Tbl. 6.8].

There are two pieces of information to gain by these tests: the resistance of the wires and whether or not they will heat too much when the currents we expect to work with are passed through them. The methods differ slightly between the on-chip wires and the "Z"-wire because the resistance of the "Z"-wire is so much lower. These will be explained in their own subsections. A second set of tests were conducted four months later by Thomas Badr. His results do not differ significantly from my own. The lack of significant change in the results is a good indication that neither my tests not intermediate manipulations had damaged the chip. Fuller data and technical details can be found in Appendix B in sections that follow the same structure as this section.

The voltage data was gathered with a Tektronix oscilloscope.[6] Voltages used for calculating resistance are averages of fifty data values, corresponding to 200 ms in the case of the ten pulse series and 100 ms in the case of the extended pulses.

3.3.1 On-chip wires

The on-chip wires had high enough resistance to be measured with a Keithley[7] multi-meter in four-wire mode. This measurement mode discounts the resistance of the wire leads between the meter and whatever they are connected to, the break-out box connected to the 32-pin multisignal connector in our case. The meter measures $\approx 5\,\mathrm{m}\Omega$ when the leads are connected so should not be trusted with measurements close to that. The results of these measurements, denoted Ω_4, are shown in Table 3.1.

To test for heating current was applied across the wire in two patterns: a series of ten one-second, high-current pulses nine seconds apart and two three-second pulses, one at low

6. Model TDS2014C.
7. Model 2110 $5\frac{1}{2}$.

Resistances of on-chip wire circuits

Wire circuit	DC50	DC100	U1	U2	U3	U4	RF1	RF2	RF3	RF4
$\boldsymbol{\Omega_4}$ (Ω)	7.37	5.94	5.88	5.66	5.99	5.33	18.13	17.65	17.53	17.84
$\boldsymbol{\Omega_{hi}}$ (Ω)	8.2	6.47	6.73	6.35	6.53	6.0	20.4	19.7	19.4	20.1

Table 3.1 — Resistances of the on-chip wires as measured in four-wire mode (Ω_4) and calculated from the voltage at the end of a three-second high-current pulse (Ω_{hi}). "High-current" is 0.5 A for RF wires and 1.5 A for the others.

current and one at high current. What constitutes "high" or 'low" current has a different meaning for the DC and "U"-wires (1.5 A) than for the thinner RF wires (0.5 A). For the wider trapping and "U"-wires the single pulses were at 0.5 A and 1.5 A while for the narrower RF wires they were at 0.15 A and 0.5 A. With the series of ten short pulses, the peaks should all rise to the same height. This indicates that the heat is being dissipated quickly enough while the current is off to return it to its original temperature after 9 s. For the single pulses the voltage needed for a given current should stay constant over the pulse.

The series of ten 1 s pulses showed no rise in voltage. Neither did the lower-current extended pulse of 3 s. The higher-current extended pulses did show an increase in resistance. The resistance reached at the end of the high-current pulse is listed as Ω_{hi} in Table 3.1. For the larger wires this was on average $(0.75 \pm 0.15)\,\Omega$ and $(2.11 \pm 0.19)\,\Omega$ for the RF wires, at least in the low-frequency limit. If all of this change in resistance is the heating of the gold wires, this translates by equation (3.5) to an increase in temperature of $(36 \pm 6)\,°\mathrm{C}$ which was not significantly different between the RF and other wires. This indicates that the wires cannot sustain these kinds of currents for extended periods but the lack of shorts or changes in the resistance after such tests indicates that this was not enough to damage them.

3.3.2 Coplanar microwave guide

The calculations and simulations to model the CPW carried out by R. Dubessy based on data collected by T. Badr with a Rohde & Schwarz[8] vector network analyzer. It is the result of his modeling that I present in 3.7.

Sodium has a mw-induced Feshbach resonance with width of 1.4 kHz when the strength of the magnetic part of the microwave field, $B_{mw} = 1\,\mathrm{G}$. This width scales like B_{mw}^2 (see Table 1.1 from [43]). Therefore, our calculations suggest that 1 A through the central strip of the CPW will give a resonance of about 100 Hz wide at the location of our trap which should not be too narrow to find.

3.3.3 "Z"-wire

The resistance of the "Z"-wire was too low to reliably measure with the multimeter so all of the resistances must be deduced from current and voltage data. The current was supplied by a Delta Elekronika power supply [9] driven by a Keysight[10] function generator via it's I_{prog} port. Current readings were taken from the power supply itself. Voltage was monitored on the $V_{monitor}$ port of the power supply. This signal was noisy so I used a low-pass filter with

8. Model ZNB20.
9. Model SM 15-100.
10. Model 33220A.

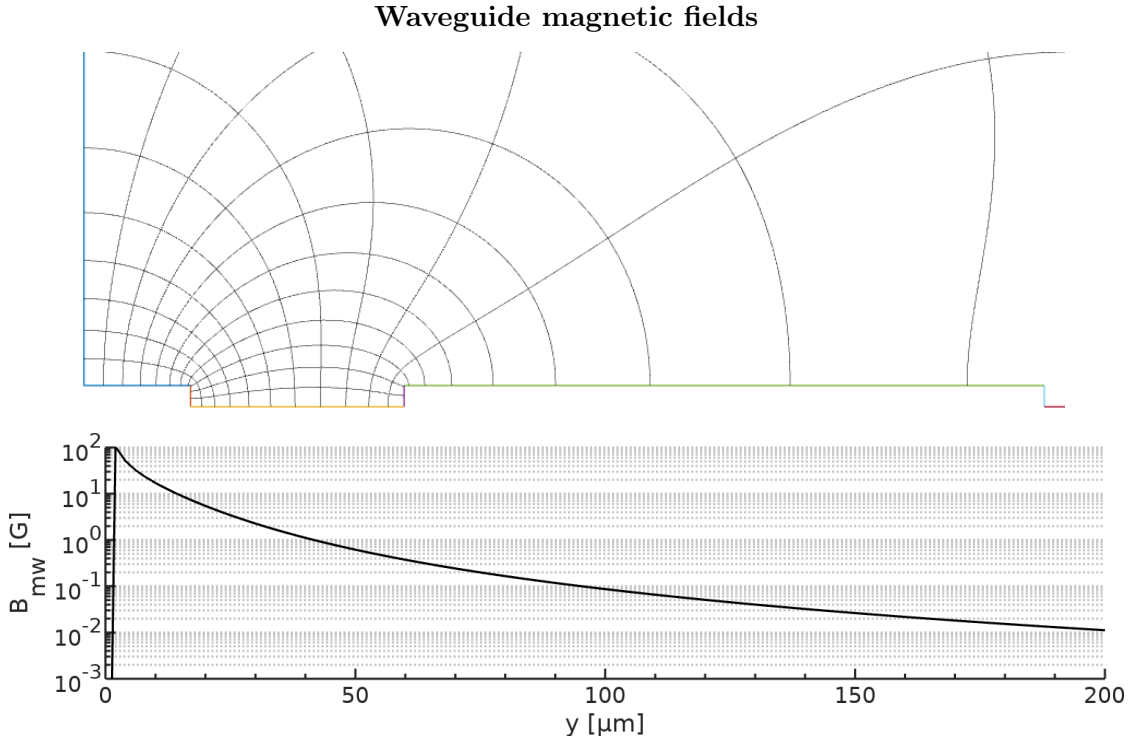

Figure 3.7 — Top — Plot of the field lines and surfaces of equipotential in half of the coplanar waveguide. Left-hand edge corresponds to the y-axis in Figure 3.2 which is a plane of symmetry. This was achieved with the Schwarz-Christoffel mapping of the complex plane to the rectangular forms representing the profile of the waveguide wires. Bottom — magnitude of the magnetic field a distance y above the coplanar waveguide. Theoretical prediction for the strength of the magnetic field a distance y above the microwave guide, carrying a current of 1 A. Courtesy of R. Dubessy.

a cutoff frequency of 16 Hz. T. Badr used a low-pass filter with $f_c = 700\,\text{Hz}$, but our results are comparable giving 5.2 to 5.7 mΩ.

One test consisted of a series of ten, 2 s, 50 A pulses 8 s apart. The idea is the same as the pulse test in the chip wires: to see if the resistance rises over these pulses. The difference between the resistances calculated by averaging over the first and last peaks comes to a negligible 5.9 µΩ, less than the standard deviation within one of the peaks, $\simeq 10\,\mu\Omega$.

The other consisted of single, extended 4 s pulses ranging in current from 25 to 100 A in 25 A increments and another one at maximum current, 112.4 A. At currents of 75 A and below, the resistance did not rise significantly. Using the formula for Joule heating, $P = I^2R$, where P is the power that goes into heating, the heat dissipation rate is somewhere close to 30 W. That assumes that all the heat generated $I = 75\,\text{A}$ is dissipated. At 100 A and above the resistance rose by a noticeable amount $\simeq 10\,\mu\Omega$. A linear regression of the data from the oscilloscope give a slope of $\simeq 1.8 \times 10^{-6}\,\Omega \cdot \text{s}^{-1}$ implying a temperature increase of $0.08\,^\circ\text{C} \cdot \text{s}^{-1}$ at 100 A or $2.5 \times 10^{-6}\,\Omega \cdot \text{s}^{-1}$ and $0.12\,^\circ\text{C} \cdot \text{s}^{-1}$ at 112.4 A. While this suggests that the heat dissipation is adequate for our expected working currents of around 80 A, it was noted that at 100 A and above, 16 s was not long enough to cool the piece down to its original resistance. When the 20 s, 20% duty cycle was ran through multiple periods, the initial voltage measured would increase very slightly ($< 1\,\text{mV}$) with each pulse.

3.3.4 Evaporation coil

The evaporation antenna itself is a Kapton^{TM}-coated 24 gauge copper wire wound in a single layer of ten coils the recess cut into the upper tower measuring 5 mm across. The center of this coil is (18.9 ± 0.1) mm below where the atoms will be held in transport magnetic trap. Two dummy coils were made to perform the tests: one with a core of wood an one with a core of copper to see what effect this will have on the magnetic field strength. The probe consisted of a small loop of copper wire and its transfer function was measured at different frequencies. Then it was held at different heights above the coil, which was driven at a frequency ν_{rf}, and the probe voltage V_p measured. The field strength is then

$$B_{rf}(\nu_{rf}) = \frac{\eta V_p}{\nu_{rf} A_p} \tag{3.6}$$

where η a factor given by the transfer function and A_p is the area of the probe.

The result is that for 0.5 W of power we expect $B_{rf} \lesssim 10$ mG in the region of our trap in the experiment at the relevant frequencies. With a wooden core, the data suggest that desired field strength of ~ 100 mG could be achieved all the way up to $\simeq 15$ MHz at 0.5 W of power. But the difference made by the copper core in the relevant frequencies is a factor of between around ten at 20 MHz to around fifty near 7 MHz. The field strength is inversely proportional to the distance from the coil and the field strength goes as the square-root of the power, as expected. But even 8 W of power was not enough to supply the desired field strength with a copper core. In the event that we cannot get enough power to the trap this way, we will have to construct external antennae.

3.4 Conclusions

The on-chip wires performed well under conditions that were harsher than they will be when they are used during the experiment. We will be able to maintain a trap with $\omega_\perp/2\pi = 2$ kHz by running 1 A through the trapping wires for at least a second without adverse heating effects. This will put the trap 100 µm above the chip which should be far enough to avoid roughness in the wires affecting the potential (Cf. [207] where the potential is relatively smooth at 107 µm). The "U"-wires proved to be stable under similar currents and will be able to provide fields of similar strength. The rf-wires will be able to carry 0.5 A of current for one second, the same time as the trapping wires and so should suffice to dress the trap (Cf. [146] where the rf-wires carried 0.6 A for times on the order of milliseconds). The CPW should provide a strong enough field to make the resonance wide enough to find.

The heat dissipation for the high-current "Z"-wire also proved adequate, at least on short time-scales. A heat-dissipation rate of 30 W is inconsistent with the observation that the voltage did not return to its original value on successive pulses of 100 A or more as the heat from such a pulse would be dissipated in under 8 s. If the flatness of the voltage during the 75 A pulse is spurious (e.g. the rise could be hidden by noise) but the one at 50 A is not, then we can only be sure from the data that the rate of heat dissipation is just under 14 W. Then the heat generated in 4 s of 100 A of current would take almost 16 s to dissipate, in keeping with observation. The "Z"-wire is expected to carry $\simeq 80$ A for 1 s every 30 s. This pulse generates 35 J of energy so even 14 W of heat dissipation is sufficient.

The evaporation coil is the only part tested that presents difficulties. The presence of OFHC copper as a core for the antenna significantly reduces its power. If this results in

excessively long evaporation times, even at high rf-power, other solutions, such as external rf-antenna must be sought.

The starting point parameters for the trap will be to run $I = 1\,\mathrm{A}$ with a transverse bias field of magnitude $B_b = 20\,\mathrm{G}$ and a longitudinal trapping bias field of magnitude $B_{ip} = 3\,\mathrm{G}$. This will make an elongated harmonic trap with a transverse trapping frequency of $\omega_\perp/2\pi = 2\,\mathrm{kHz}$ and an aspect ratio of $\omega_\perp/\omega_x \gtrsim 100$. This will sit around $100\,\mu\mathrm{m}$ from the chip. At this distance 1 A of current in the CPW will give the magnetic field component of the microwave field a magnitude of just under 100 mG (see lower register of Figure 3.7). If this makes the resonance too narrow (recall $\Delta\omega \propto B_0^2$) the first thing to try is to lower the current in the trapping wire and move the trap closer to the chip since the etching process has already reduced the width of the CPW microstrip by almost half its designed width. We do not have quantitative data about the roughness of our wires which would allow us to estimate how it would effect the trap at different distances [221], but there are tricks relying on alternating current to deal with the effect of wire roughness on the trapped cloud [224] that could let us get us get within $\sim 10\,\mu\mathrm{m}$ of the chip surface, and get as much microwave power as we need.

The chip has everything we need to trap and manipulate the atoms and look for the Feshbach resonance that will be the first goal of the experiment. The second trapping wire and the four rf-wires interleaved between the other wires on the chip should give the experiment a long lifetime. These give us options to explore other physics such as double-wells and adiabatic rf-fressed traps. This is before we add any optical equipment and try to make a homogeneous box potential for direct comparison to the Lieb–Liniger model. And we should not need to break vacuum if we have chosen our materials and constructed the parts well enough.

Chapter 4

Objective lens design

This chapter is devoted to the design of three objective lenses that will be used in the science chamber. It is best thought of as a pair of lenses: a low-magnification objective and a high-magnification objective that can be converted for in-situ imaging. The low-magnification lens will be for imaging the cloud in, or released from, the transport magnetic trap in the science chamber, before it has been loaded into the trap on the atom chip. The high magnification objective will be needed for time-of-flight (ToF) imaging after the gas has been released from the atom chip to measure density fluctuations that arise from phase fluctuations in the trapped gas [184]. The conversion to in-situ imaging is useful, but this is ancillary to the ToF imaging which will give us our most important data in the early phases of the experiment. The planned imaging setup is a typical absorption imaging one [31, §3.2]. I took for my starting point set of lenses described in [225, §2.1.3] and I will be taking a few formulae from his appendix on optics and imaging. For foundational optics, including geometrical optics, see [226].

The problem can be conceived as an optimization problem on a function of many variables. One wants as much resolution over as much of three dimensional space as possible. But, as a general rule, one must trade off one thing for another. A simple example would be that if the magnification is increased, less area in object space can be imaged. The foundation of this design work was finding a region in this parameter space where all the performance characteristics where acceptable.

I designed the lenses in the Zemax lens design software to be constructed using commercially available lenses and optomechanical parts. The only exception being the apertures for the high-magnification and in-situ lenses which can be machined from simple metal blanks. This saves the cost of custom-made parts and lenses. Another simplifying feature is that the lenses are *infinite systems* (or nearly so) meaning they collimate the light at some point and they can be thought of as having a front collimating group and a rear refocusing group. This means that the distance between these groups is arbitrary, to the extent the beams are columnated. This provides a little ease of construction in that the lengths between the groups, as well as the location of the apertures, need not be constructed to precise tolerances but care would need to be taken in the focusing if they were placed very far apart (e.g. using a mirror and having a track length of meters as in [225]). It also gives some conceptual simplicity in designing them as it means the overall magnification of the lens will just be the ratio of the effective focal length of refocusing group to the effective focal length of the collimating group.

In Section 4.1 I go over the background needed to understand the design decisions and

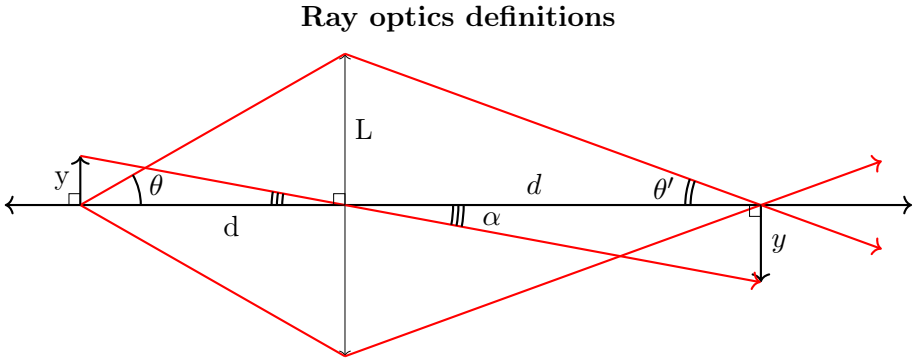

Figure 4.1 — A *simple* (i.e. single) magnifying lens of radius L with two different numerical apertures: $\mathrm{NA}_{obj} = \sin\theta$ and $\mathrm{NA}_{img} = \sin\theta'$. Rays leaving a point in the object plane a distance d from the lens in object space will meet at the image plane a distance d from the lens in image space. The size of the image, y, is determined by the chief ray, traced from the edge of the object, a distance y from the optic axis, through the center of the lens to the image plane. Since this ray passes through the lens without deviation it forms an angle α with the optic axis in both object and image space.

simulation data. In Sections 4.2 and 4.3 I detail the two lenses and give their expected performance. The conversion of the high-magnification objective for in-situ imaging will be covered in the section covering the high-magnification lens. Finally, in 4.4 I draw my conclusions. Performance charts are at the end of the chapter. Lens schematics including part numbers can be found in Appendix C.

4.1 Background

First I will go over the physical considerations such as the constraints of the science chamber. Then I will give the definitions and formulas necessary to design and understand these lenses.

4.1.1 Physical parameters

The science chamber is a stainless steel octagon about 17 cm across with circular windows of diameter 35.1 mm at the air-side flange. The flange is held in place by bolts with heads 5 mm thick whose centers are around 60 mm apart across the window. I add an extra 3 mm to the thickness of the bolt heads to let the retaining rings, of thickness around 2.5 mm, sit in the lens tube. The atoms are held in the center of the chamber 81.8 mm away from the inside of the window glass. This gives the window passing light from the atoms a numerical aperture, NA, (see next section) of 0.2. The atom chip is suspended from the top of the chamber with its surface about 3 mm above the center of the window. Our camera is a Pixelfly[1] comprising a 1392×1040 sensor of 6.45 µm square pixels making it roughly $9 \times 6.7\,\mathrm{mm}^2$ in size. A lens whose housing is less than 50 cm long is preferable as the shorter the overall length of the lens, the less susceptibility to vibration and the easier to fit on the table.

4.1.2 Optical formulae

Here I lay out definitions and concepts I will use throughout this chapter. Throughout I will use unprimed symbols in roman type to indicate quantities in *object space* (the side of lens where the object is) and primed or italic symbols to indicate quantities in *image space* (the side of the lens where the sensor is).

First is the numerical aperture, NA, which is the sine of the half the angle made by the cone of light either emitted from a point on the optic axis and entering on the object side (NA_{obj}) of a lens or leaving it and being focused to a point on the optic axis on the image side (NA_{img}.) This will, in general, depend on where on the optic axis these points falls and not just on the lens. The NA gives us a convenient way of defining magnification, M, and translating between lengths in image space and object space. Magnification is the ratio of the size of an image to the size of its object, so $M \equiv \frac{y}{\mathrm{y}}$ in Figure 4.1. Since the angle made by the chief ray to the optic axis, α, is the same on both sides of the lens $\tan\alpha = \frac{y}{d} = \frac{\mathrm{y}}{\mathrm{d}}$. And since the radius of the lens, L, is the same on both sides of the lens

$$\frac{\tan\theta}{\tan\theta'} = \frac{\cancel{L}}{\mathrm{d}}\frac{d}{\cancel{L}}. \tag{4.1}$$

In the paraxial approximation, the angles θ and θ' are assumed to be small so $\tan\theta \approx \sin\theta$. Indeed, our largest NA is .2 meaning our largest $\theta = \arcsin(.2) = .2014$ and $\tan\theta = .2041$. Putting all this all together we get $M \equiv \frac{y}{\mathrm{y}} = \frac{d}{\mathrm{d}} = \frac{\tan\theta}{\tan\theta'} \approx \frac{\sin\theta}{\sin\theta'}$ and recalling $\mathrm{NA} \equiv \sin\theta$ allows the magnification to be defined in terms of the numerical aperture

$$M \equiv \frac{\mathrm{NA}_{obj}}{\mathrm{NA}_{img}}\ . \tag{4.2}$$

This definition also saves us from having to put absolute value marks around every instance of M but it should be noted that these lenses are inverting and so have negative magnification when defined as the ratio of displacement in image space to displacement in object space.

A few words on the modulation transfer function (MTF) will also be useful for understanding this chapter. If one images a series of black and white bars of various spatial frequencies one will not get a perfect reproduction of pure black and pure white bars. Obviously, at some high frequency the bars will be unsolvable, but even before that the edges will blur and the bars will appear more and more gray as the frequency rises. The average of the ratio between the "black" of black bar and the "white" of the white bar next to it is called the contrast ratio and when this is plotted against spatial frequency one obtains the MTF. Where this reaches zero is called the cutoff frequency [225, A.21],

$$\nu_c = 2\frac{\mathrm{NA}}{\lambda} \tag{4.3}$$

for incoherent light,[2] and is one way of discussing the smallest resolvable object. To get an idea how the lens behaves off-focus one can define a *through-focus MTF* where a particular spatial frequency is picked and the contrast ratio is plotted against the object's position as it is moved along the optic axis. This is a way of looking at the how things fall out of focus on-axis for a particular spatial frequency or feature size.

Zemax has some idiosyncrasies to how it plots its modulation transfer functions such as those presented here. The MTFs used here are called fast Fourier transform (FFT) MTF

1. Model PF-M-QE from PCO
2. for coherent light it is half this number.

because Zemax calculates them by taking the FFT of the system's *pupil function*, the function which describes the image the system forms of the stop aperture. The horizontal axis of the MTF graphs Zemax produces (see Figure 4.5 and Figure 4.6 (d.)) are given in inverse millimeters *in image space*. If one wants to know what spatial frequency this corresponds to in *object space* one must multiply by the magnification of the lens. This is all the knowledge of the MTF I will be calling on here. For more details see [225, §A.5.1] and references therein.

A key piece of data output by Zemax is the *spot diagram*. To make a spot diagram one (Zemax in the present case) traces some set rays originating at a point and traveling through the stop aperture of the lens; making sure to include the *marginal rays* (i.e. the rays passing the edge of the aperture). The diagram then consists of all the points where a ray lands on the image surface. This captures all of the geometric or spatial (as opposed to phase or wavefront) aberrations the lens introduces to the light path. When done in Zemax there is an option to plot the Airy disk formed by the lens as a black circle. If all the points on the spot diagram fall within this circle, the aberrations can be more or less ignored as all the points originating from a point are falling within the smallest circle that point could be imaged as. One can do this for different points in the object plane, which Zemax calls "fields" and enumerates beginning with "field 1," to get a sense of how the lens behaves off-axis.

I will be using several image quality parameters such as resolution and depth of field (DoF) in order to evaluate my designs. In the paraxial approximation, if rays from a point in the focal plane form a point in the image plane, basic geometry shows that the rays from a point located a distance, dx, away from the lens, along the optic axis will form a circle in image space of diameter, dz, whose formula would be written

$$\begin{aligned} dz &= 2\sin\theta\,\mathrm{dx}\,M \\ \implies \mathrm{dx} &= \frac{dz}{2M\,\mathrm{NA}_{obj}}\ . \end{aligned} \tag{4.4}$$

It is easy to see why this should be NA_{obj} for a simple lens but in that case it makes no difference as there is no magnification and so both numerical apertures are equal. But take the case of an infinite system: it is the objective NA which determines how far from collimation an object's rays will deviate and thus how far from being focused by the refocusing group focus it will be. If dz is taken to be the maximum acceptable spot size, the pixel size, or *pixel pitch*, in image space, l_{px}, for instance, then dx becomes the minimum *geometric* depth of field, DoF_g, on one side of the plane of best focus. Doubling this gives what I will use as the entire geometric DoF,

$$\mathrm{DoF}_g = \frac{l_{px}}{M\,\mathrm{NA}_{obj}}\ . \tag{4.5}$$

This is in fact a little pessimistic. This insists that the circle of confusion be split evenly between two pixels in the worst case. This would be (naively) expected to give an MTF/contrast ratio of 0.5, still perfectly legible. So one can expect the actual depth of field to extend some ways beyond this.

If one instead imagines a point in object space being imaged as a spot at the center of a pixel and moves that point object in the focal plane off the optic axis by an amount dy, this will move its image in the image plane by an amount Mdy. If one insists that the minimum acceptable change in image space is a pixel one obtains the geometrically limited resolution

$$\mathrm{dy}_{min}(M) = \frac{l_{px}}{M}\ . \tag{4.6}$$

This is only an estimate. One could argue, just for example, for a minimum change in image space of $\frac{\sqrt{2}}{2}l_{px}$ as a point is, at worst, imaged in the center of a pixel and further than $\frac{\sqrt{2}}{2}l_{px}$ would put the new image on a pixel diagonally adjacent; again the worst case scenario. Of course lenses do not form point-like images even of infinitesimal objects, but Airy disks, so this formula is more useful in conjunction with (4.5) for getting a general estimate of whether diffraction or geometric effects limit the resolution.

The formula for the angular diameter of an Airy disk formed by light of wavelength λ passing through a circular aperture,

$$\alpha_A = \frac{1.22\lambda}{D}, \tag{4.7}$$

is well known. Here α_A is the angle subtended by the ring of the first minimum of the Airy disk and D is the aperture diameter. Since numerical aperture can be defined in terms of the angle marginal rays focused from infinity make with optic axis we may divide this by 2 and construct the Airy radius in terms of the NA by taking $\tan\theta \approx \sin\theta$ for small θ and $\sin\theta = \mathrm{NA}_{img}$ since the Airy disk forms on the image side of the lens. In air one obtains

$$r_A = \frac{0.61\lambda}{\mathrm{NA}_{img}}\,. \tag{4.8}$$

Not that NA in equation (4.8) must be NA_{img} *not* NA_{obj} as is often stated in the literature on microscopes. If one takes the Rayleigh criterion for resolution, two resolved airy disks are separated by this r_A, and this becomes the smallest resolvable image. To find the smallest resolvable object one simply takes $\mathrm{NA}_{img} \to \mathrm{NA}_{obj}$. This gives the formula for *diffraction limited* resolution the form

$$\mathrm{dy}_{min}(\mathrm{NA}) = \frac{0.61\lambda}{\mathrm{NA}_{obj}}\,. \tag{4.9}$$

The MTF cutoff frequency will imply a smallest resolvable object just slightly smaller than this due to the contrast ratio being slightly above zero at the Rayleigh condition.

Off-focus, the lens will not form a perfect, diffraction limited Airy disk but will lose resolution in non-trivial ways and one can define various criteria for being no longer resolved. I take R. Bücker's formula [225, Eq. A.24] for the *diffractive* depth of field, DoF_d, where things fall out of focus due to diffractive effects with one modification. R. Bücker gives the distance from the plane of best focus where detail is lost, I double this to get the total depth of an object that can be held in focus. This gives

$$\mathrm{DoF}_d(\mathrm{NA}) = \frac{\lambda}{\mathrm{NA}_{obj}^2}\,. \tag{4.10}$$

These figures of merit which indicate whether the DoF and resolution are limited by diffractive or geometric effects are plotted in Figure 4.2.

Since designing a lens is a constant compromise, it is nice to have ways to put limits on quantities, beyond which they are not useful and can be given up for something else. A limiting balance can be precisely constructed between the magnification and the numerical aperture using the Nyquist-Shannon sampling theorem [227]. This puts a limit on the spatial frequency which can be uniquely reconstructed, called the *Nyquist limit*, $\nu_N = \nu_s/2$, given a sampling frequency ν_s. This limit can be used to determine how much magnification is needed to uniquely reconstruct a signal appearing in object space, any more being redundant. If one imagines the diffraction limited image as a series (or grid but our problem is symmetrical so

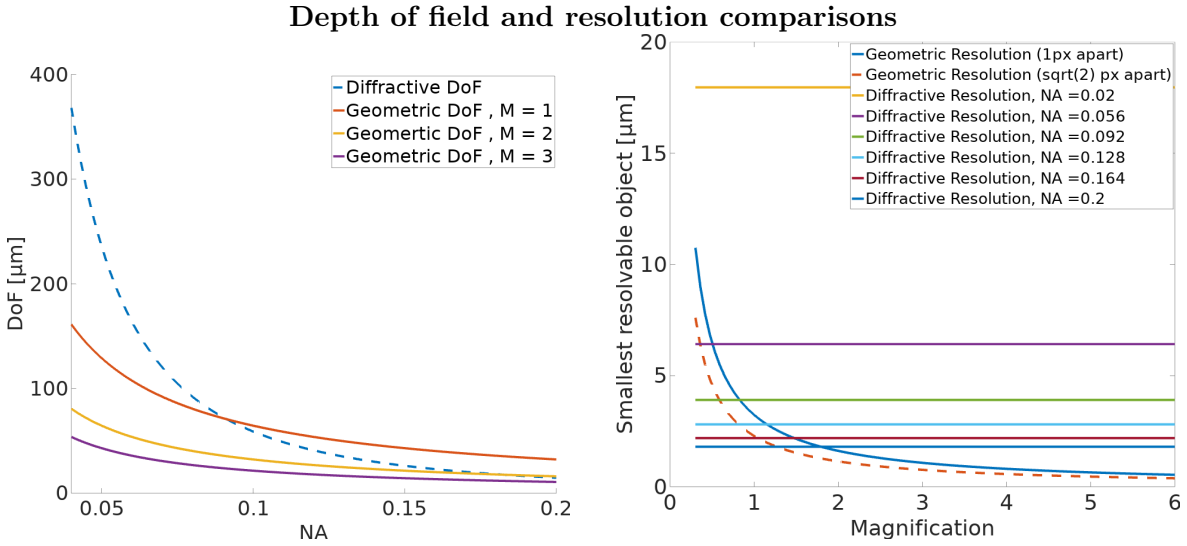

Figure 4.2 — Plots of the DoF equations (4.5) and (4.10) and the resolution equations (4.6) and (4.9) for various values of NA and M. These are not hard limits and are most useful for finding what is limiting the DoF or resolution. Left — Take the smaller value at any point to determine whether geometrical or diffractive effects limit the DoF. Note that geometrical effects limit the DoF for all $M \gtrapprox 2$, NA ≤ 0.2. Right — Take the larger of the smallest resolvable objects at any point to determine whether geometrical or diffractive effects limit the resolution.

I take only one dimension) of Airy disks separated by r_A, this will give an image with spatial frequency of r_A^{-1}. The optimal magnification will put two pixels inside each of these cycles of one Airy radius. Then one can substitute equation (4.2) into equation (4.8) multiplied by two to obtain a formula for the ideal magnification for a given numerical aperture. For a sensor with pixels l_{px} square this formula reads

$$M_{\max} = \frac{2\,\mathrm{NA}_{obj} l_{px}}{0.61\lambda}. \tag{4.11}$$

Magnification beyond this limit can be called *empty magnification* and can be traded off for *either* more diffraction limited resolution *or* DoF but, by the equations above, not both.

To treat questions related to the field of view (FoV) as well as DoF, one needs some criterion for what needs to be kept in frame or in focus. Let the *four sigma edge point* or just *edge* of a cloud be the point in the density distribution two standard deviations away from the center, the location of maximum density. If one can image both of these edges one can see upwards of 95% of the atoms and should be able to extract a reliable fit of the profile.

For discussions and explanations on kinds of aberrations, coma, spherical, astigmatism and so forth, see [225, §A.2] or [226, §4.5ff.]. The distinctions are not critical to this thesis as long as one accepts that aberrations are undesirable and cannot be ignored when they cause rays from a point to fall outside of the lens' Airy disk and thus a affect the image.

4.2 Low magnification

The low magnification objective will be used to image the "hot" ($\lesssim 200\,\mu\mathrm{K}$) cloud in, or released from, a magnetic trap centered $\approx 3\,\mathrm{mm}$ below the surface of the atom chip. It will

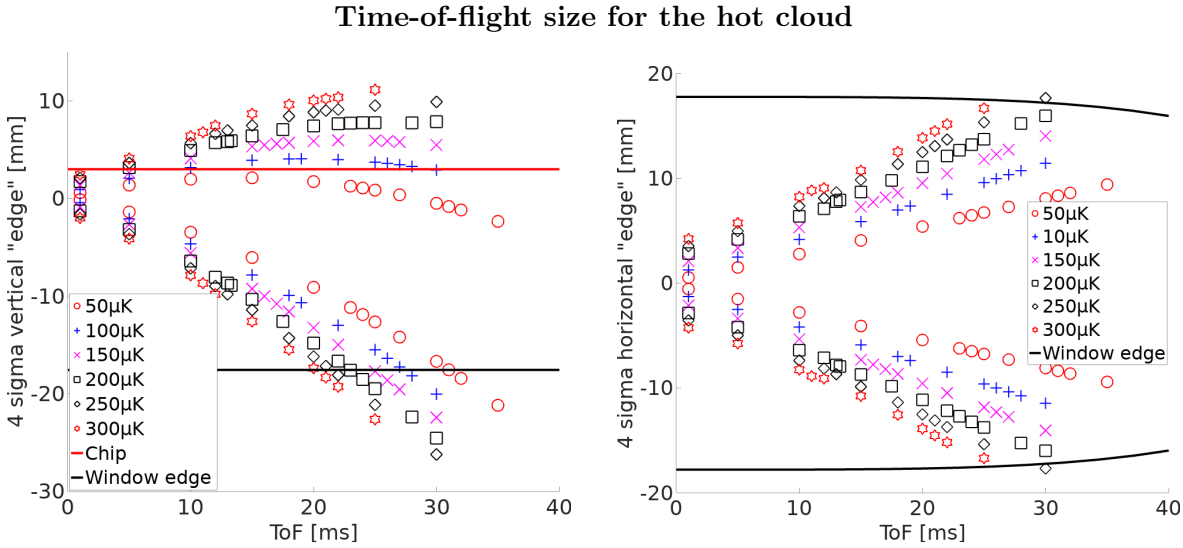

Figure 4.3 — Dispersion of a "hot" cloud at various temperatures after time of flight in the horizontal and vertical directions in order to determine the required field of view. Here the cloud edge is defined as the point in the density distribution two standard deviations (calculated with (4.13) and its vertical analog) away from the cloud center. If both these points are visible, enough of the distribution should be visible to perform meaningful curve fitting. Left — Vertical position of the top and bottom edges of the cloud after time of flight along with the bottom of the window and the chip surface. Right — Horizontal position of the left and right edges of the cloud after time of flight along with the window edges which curve inward towards the cloud as it falls.

be used for absorption imaging in order to see the general shape and location of the cloud and to extract an atom number/density profile. The cloud itself will be on the around two centimeters in size on its long axis (see Figure 4.3) so resolution on the order of a millimeter will suffice. Extracting an overall atom number only relies on imaging all the light that is absorbed by the cloud and extracting a temperature requires only obtaining variances of the expanding density profiles and fitting with equation (4.13) or fitting the density profile of a single image (see Appendix E). In any case, resolution is not critical and need only be a little below a millimeter. Because resolution is not critical, single lenses or lens groups (i.e. doublets etc.) should suffice for the front and rear groups.

The more important consideration for this lens is that the entire cloud is imaged; the wings of the distribution being most critical for fitting with a unique Gaussian. If one has N thermalized atoms, each of mass m, in a linear trapping potential $V(\boldsymbol{r})$, one may write the density at temperature T as

$$n(\boldsymbol{r}) = C \exp\left(-\frac{V(\boldsymbol{r})}{k_B T}\right), \tag{4.12}$$

where C is a normalization constant such that $\int n(\boldsymbol{r}) d\boldsymbol{r} = N$ and k_B is Boltzmann's constant. See § 5.1.1 for details about the potential. The density distribution after ToF, t, will be this density distribution convolved with a Boltzmann velocity distribution to model the free expansion. The collision rate of the cloud in the magnetic trap is around once per second (see 5.2.2) so it can be assumed that the atoms do not interact in the few milliseconds after they are released from the trap. Therefore the variance along the x-axis after t separates into two

Low-magnification objective

Surface	Radius	Thickness	Glass	Semi-Dia.
1 Object	∞	81.8	vacuum	—
2 Window glass	∞	6	Silica	18.35
3 Window flange	∞	6.5		17.55
4 Bolt heads	∞	8		30
5 Air gap	∞	196.411		—
6 300 mm doublet	557.4	2	SF2	12.7
7	137.09	4	N-BK7	12.7
8	-165.2	(20)		12.7
9 Stop	∞	(20)		(3.00)
10 100 mm doublet	62.75	4	N-BK7	12.7
11	-45.71	2.5	SF5	12.7
12	-128.23	79.535		12.7
13 Camera opening	∞	17.526		12.7
14 Sensor	∞	—		—

Table 4.1 — Surface specifications for the low-magnification objective. All units are millimeters. Lengths in parentheses are variable.

parts, reading

$$\left\langle (x - x_0)^2 \right\rangle = \frac{N}{\alpha^2} + N\frac{k_B T}{m} t^2, \tag{4.13}$$

where $\alpha = g_F\, m_F\, \mu_B\, b'/(k_B T)$ with, in our case, g_F the Landé factor of $-1/2$, m_F the total spin projection on the quantization axis of -1, μ_B the Bohr magneton and b' the horizontal magnetic field gradient of $65\,\mathrm{G}\cdot\mathrm{cm}^{-1}$. In the vertical direction, the magnetic gradient is twice as large so the first factor on the right-hand side, accounting for the initial distribution, will be divided by four. There will also be a factor of $g^2/4$ in the last term for the acceleration due to gravity.

Equation (4.13) allows us to find the four sigma edge point of the cloud and plot the evolution of these end points at various temperatures in Figure 4.3. From this we can see that the camera is best oriented landscape-wise with the long axis horizontal and that the limiting distance will be in the vertical direction and correspond to a length of around 20 mm. Given that the short axis of our sensor is 6.7 mm this gives a desired magnification $M = .335$. This would give a horizontal FoV of ≈ 27 mm, corresponding to an ample ToF of 20 ms for even the highest temperature clouds plotted in 4.3. A typical cloud in our magnetic trap is 185 µK, just about in the middle of the range of temperatures plotted. For comparison typical ToF, such as in Figure 2.5, is around 6 ms.

4.2.1 Lens design and performance

The compromise when choosing lenses for the low magnification system is between choosing lenses so long that the whole setup is too far away and choosing lenses so short that they introduce aberrations and vignetting. The choice presented here of an $f = 300$ mm achromatic doublet for a collimating lens represents the shortest, readily available choice which allows for a refocusing lens that is both short enough to give the desired magnification, $f = 100$ mm $\Longrightarrow$

$M \approx .333$, without also introducing too much aberration at a convenient lens diameter, $D = 1" \approx 25.4\,\text{mm}$, the diameter of the C-mount lens aperture of our camera.

With this arrangement and a stop aperture of $D = 12\,\text{mm}$, the largest allowed by an in-tube diaphragm of the correct tube size, showed significant astigmatism away from the optic axis with this setup. Image simulation in Zemax indicated a reduction in image quality at the corners of the sensor. This implies that it would be ill-advised to move to a shorter lens combination or a wider aperture as these would only worsen the aberrations. Choosing longer lenses is inconvenient as the next 1:3 ratio up in ThorLabs' catalog would be an $f = 250\,\text{mm}$, $D = 1"$ imaging lens and an $f = 750\,\text{mm}$, $D = 2"$ collimating lens. But since a lower resolution causes far less difficulty, I just reduced the aperture opening to 6 mm, taking NA_{obj} from .02 to .01, and expect the results shown in Figure 4.5. There remains some astigmatism outside of the Airy disk in the very corner of the sensor (field 4,) but not enough to change the cutoff of the MTF and everywhere else is diffraction limited.

By comparing equations (4.5) and (4.10) one can see that geometric effects still dominate the DoF, given our demanding circle of confusion criterion. But even on the condition that the circle of confusion fit on a pixel this lens will still give a $\text{DoF}_g \geq 1.9\,\text{mm}$. This reduction in aperture does result in some loss of resolution, changing in the cutoff frequency of the MTF $\nu_c \approx 200\,\text{mm}^{-1} \rightarrow \nu_c' \approx 100\,\text{mm}^{-1}$ which corresponds to a doubling of the size of the smallest resolvable object from $\approx 15\,\mu\text{m}$ to $\approx 30\,\mu\text{m}$. This is in close agreement with equation (4.9), which gives 36 µm, and not exactly catastrophic.

But even taking this absolute resolution limit to be a bit optimistic the performance is fine. Taking a 50% contrast ratio as being where an object will be clearly imaged, one can see from Figure 4.5 a) that for an aperture of 6 mm this corresponds to $\approx 40\,\text{mm}^{-1}$ in image space or an object size of $\approx 75\,\mu\text{m}$; plenty of detail. A variable iris aperture allows some flexibility: either stopping down to gain more DoF at the cost of resolution for a larger cloud, or opening up to gain more resolution on a smaller cloud at the cost of aberrations on the corners where the cloud is less likely to be anyway. Although not completely clear from Figure 4.2 (b), the resolution is in fact diffraction limited as equation (4.5) gives a geometric resolution $\text{dy}_{min}(M) \approx 20\,\mu\text{m}$.

The final design is laid out in table 4.1 and figures (4.5) and (C.1) and can be constructed entirely from parts ordered from Thorlabs and assembled in their $D = 1"$ SM1 lens tubes (see Figure C.1 for details.) Given the tube choices from the Thorlabs catalog and the requirement of having an iris in between the lenses with enough room to spare on the front to be safe, the final design amounts to 6" or about 150 mm as seen in Figure C.1. As designed, Zemax gives the $\text{NA}_{obj} = .01$ and $|M| = .3325$ almost exactly the ideal derived at the end of the last section.

4.3 High magnification

The high magnification objective will be used to photograph along the transverse axis of expanding gas after the trap on the atom chip has been switched off and the atoms are allowed to fall and expand freely. The structures to be imaged are the density fluctuations on the order of $\ell_\phi \sim 10\,\mu\text{m}$ [88, 116]. To get clear images of them we would like a resolution of around 1 µm, but a resolution up to 3.5 µm is acceptable so long as it allows for sufficient DoF to let the atoms expand for 5 to 10 ms. This will render a feature of 10 µm at a "resolution" of three distinct objects and 10 ms is enough time to allow phase fluctuations in the trap to become density fluctuations in the cloud and to make the cloud diffuse enough to be imaged

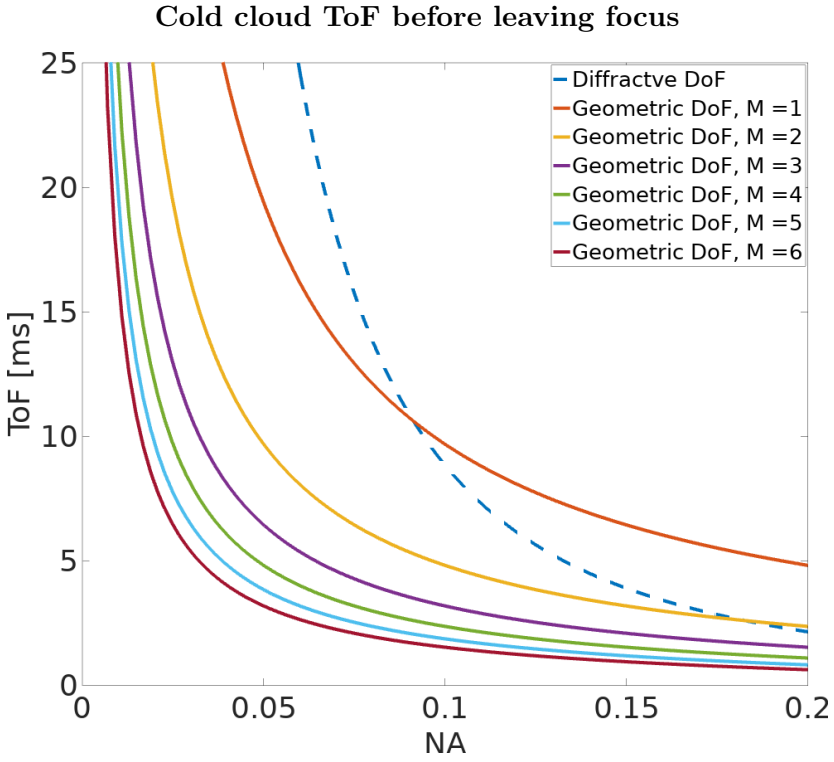

Figure 4.4 — Necessary depths of field given by different numerical apertures and magnifications needed to contain both edges of a 1-D gas released from a $\omega_\perp/2\pi = 2\,\mathrm{kHz}$ trap (right axis) and the corresponding times at which those edges fall out of focus (left axis.) Right and left vertical axes denote equation (4.15) and its inverse, respectively.

by absorption. If this does not occur, high resolution is useless.

Diffraction limited resolution, equation (4.9), is maximized by maximizing the numerical aperture which involves having the widest possible aperture and a collimating group situated as close as possible to the science chamber. Geometric resolution, equation (4.6), is achieved with a high magnification which involves having a short lens group near the object and refocusing with a longer one near the sensor. However, the longer the rear lens group, the longer the overall length of the lens and therefore the more sensitive to vibration and less convenient to fit on the optical bench. This means that I want the shortest possible front group so that the front group need not be too far away nor the rear group prohibitively long. Given the constraints of our geometry and readily available lenses the front group presented here is the shortest possible (see next section for details).

The importance of resolution does not mean one can ignore how much of the cloud we image but, for all practical purposes, the cloud will always be smaller than the image frame. And since high resolution implies a shallow DoF, the limiting factor will be the cloud falling out of focus along the optic axis rather than out of the FoV. Looking ahead to equation (4.15), the ToF for reaching the edge frame is around 14 ms for even an absurd magnification of $M = 10$. The DoF in that case would be 6.4 µm implying a ToF of under a millisecond. So it can be safely disregarded in favor of the DoF, whose requirements depend on the properties of the expanding gas.

In the trap, the confinement in the transverse axis is much stronger than in the longitudinal one, thus expansion in the transverse directions (y, vertical axis, and x, parallel to optic axis)

is expected to be much more rapid than expansion along the longitudinal axis, z. I further assume that it is in the one-dimensional regime so all of the atoms are in the ground state of the transverse trap with a trapping frequency of $\omega_\perp/2\pi = 2\,\mathrm{kHz}$. Then from [228] I take the wave function for the ground state of a harmonic trap quickly switched off and allowed to fall along the y-axis under gravity. This can be written as the ground state of a harmonic oscillator with a coordinate shift, $y \to \tilde{y} = -gt^2/2\sqrt{1+\omega_y^2 t^2}$, with gravity imparting an acceleration g and taking the initial vertical position of the trap center, y_0, to be zero. One calculates the probability density function, $|\psi|^2$, the positions of the four sigma edge points after time, t, then read

$$y(t) \pm 2\sigma_y(t) = \frac{1}{2}gt^2 \pm 2\sqrt{\frac{\hbar}{2m\omega_\perp}(1+\omega_\perp^2 t^2)}\,. \tag{4.14}$$

There is no danger of losing the edge of the cloud outside of chamber window or off the edge of the sensor so the use of equation (4.14) is to get an idea of how long a ToF we can expect that would be sure to keep all four standard deviations in focus. In the direction normal to the sensor, x, the width of the cloud should be identical to its width in the y direction but without the correction for gravity shifting the center of the distribution [228, Eq. 7.]. The ideal (or minimum if one takes the assumptions behind the DoF equations in § 4.1.2 as the most stringent) DoF for a given ToF, t, is then trivially derived from equation (4.14) and reads

$$\mathrm{DoF}(t) = 4\sqrt{\frac{\hbar}{2m\omega_x}(1+\omega_x^2 t^2)}\,. \tag{4.15}$$

This equation can also be inverted to calculate at what time the cloud will be a certain width as was done to produce Figure 4.4.

4.3.1 Lens design and performance

If we insist on a ToF between 5 and 10 ms it can be seen from equation (4.15) or Figure 4.4 that we need a DoF $\gtrapprox 35\,\mu\mathrm{m}$. If we further insist on a resolution of 3.5 µm or better we can use equation (4.6) to see that we need $M \geqslant 1.85$ or solve equation (4.9) for NA_{obj} and substitute it into equation (4.5) to find that $M \leqslant 1.795$. This would be a contradiction if the edges of the DoF were hard limits but this tells us that we cannot stray too far from a solution where $M \approx 1.8$ and, by equation (4.9), $\mathrm{NA}_{obj} \approx .1$. As per Figure 4.2 (a), or comparing equations (4.5) and (4.10), one can see that DoF will be limited by DoF_g.

To decrease the overall length of the system while getting enough magnification I use a telephoto group as discussed in [225, §A.4.6]. This involves placing a lens with a negative focal length at the rear of the system to increase the effective focal length of the refocusing group with a less-than-proportional increase in length.

Since the demands on magnification and resolution are not extreme, this could, in principle, be accomplished with just a pair of achromatic doublets. However, maximizing the objective numerical aperture will also allow for easy conversion to a high-resolution lens for in-situ imaging. Therefore I begin by placing a positive meniscus lens as close as I can to the science chamber (similar to the light-sheet imaging objective outlined in [225, §3.5.1], based on [229]) that will allow the shortest available achromat to collimate the light from the atoms. This will give the shortest effective focal length to the front group, while allowing the rear group to be as short as possible. To maximize light collection the objective lens should be large enough to cover the window in addition to being as close as possible to it. It should also have a short focal

High-magnification objective

Surface	Radius	Thickness	Glass	Semi-Dia.
1 Object	∞	81.8	Vacuum	—
2 Window glass	∞	6	Silica	18.35
3 Window flange	∞	6.5		17.55
4 Bolt heads	∞	8		30
5 Air gap	∞	1.44		25.4
6 250 mm meniscus lens	-224.67	5.53	N-BK7	25.4
7	-82.540	56.771		25.4
8 250 mm doublet	459.2	2	SF2	25.4
9	111.69	7.5	N-BK7	25.4
10	-137.09	(10)		25.4
11	∞	(10)		15
12 200 mm doublet	109.86	8.5	N-BK7	25.4
13	-93.11	2	SF2	25.4
14	-376.25	133.877		25.4
15 −300 mm meniscus lens	-95.11	3	N-BK7	12.7
16	-250	56.063		12.7
17 Camera opening	∞	17.526		12.7
18 Sensor	∞	—		—

Table 4.2 — Surface specifications for "high" magnification lens with $M = 1.8$ and ample DoF to allow a ToF of 5.1 ms. Resolution is 3.6 µm or better. Thicknesses are measured on the optic axis and all units are mm. Lengths in parentheses are variable as this is an infinite system.

length itself so as not to lengthen the effective focal length of the front group too much. Given our geometry and the scarcity of large positive meniscus lenses, the choice of a Thorlabs LE 1613-A[3] $D = 2$", $f = 250$ mm, positive meniscus lens allowed the shortest achromatic doublet having $D = 2$", $f = 250$ mm. This combination has an effective focal length, $f_e = 142$ mm when focused on the atoms. Looking at line 7 in table 4.2 one can see that there is less than 20 mm to spare between the achromat and the bolt heads and since the next shortest achromat one could hope to find would be at most $f = 225$ mm, a shorter collimating achormat cannot be chosen. An aperture of $D = 30$ mm gives an NA = .105 $\implies$ $\mathrm{dy}_{min}(\mathrm{NA}) = 3.4$ µm by equation (4.9), just enough resolution.

A doublet with $f = 200$ mm on its own as a refocusing group would give $M = 1.4$ so a fairly weak negative telephoto element is required to bring it up to $M = 1.8$. One could imagine using a slightly longer focal length achromat as a refocusing group but this extends the overall length and makes change over to in-situ imaging more difficult. I chose a lens with $D = 1$", $f = -300$ mm negative meniscus lens. At best focus, this gives $M = 1.8 \implies \mathrm{dy}_{min}(M) = 3.6$ µm and $\mathrm{DoF}_g = 34$ µm which allows a ToF= 5.1 ms. We could also squeeze out a little more resolution by refocusing the rear group.[4] For example $M = 1.85 \implies \mathrm{dy}_{min}(M) = 3.49$ µm and $\mathrm{DoF}_g = 33.2$ µm which allows a ToF= 5 ms almost exactly with no

3. "A" is the Thorlabs product for visible range (350 to 700 nm) anti-reflective coating which we use for all our lenses.
4. In table 4.2 lines 14 : 130.305 mm → 123.665 mm and 16 : 61.894 mm → 73.06 mm

appreciable affect on the MTF. But $M = 1.8$ represents the minimum of our requirements as calculated by equations (4.6), (4.5) and (4.15) and is at best focus. This design is presented in table 4.2 and Figures (4.6) and (C.2). But recall that equations (4.6) and (4.5) do not represent hard limits and are based on fairly pessimistic assumptions. If one relaxes the condition on the maximum acceptable spot size from l_{px} to $\sqrt{2}l_{px}$, then ToF $= 5\,\text{ms} \rightarrow 7\,\text{ms}$ moves comfortably into our desired range. Using an adjustable-length tube near the camera allows for such minor adjustments. As can be seen from Figure 4.6 c) aberrations do not take the spot outside of the Airy disk even at 1 mm off axis in object space. This means the cloud can expand to over two millimeters in size without significant loss of resolution due to aberrations.

Unfortunately, this lens has an Airy radius of $r_A = 6.5\,\mu\text{m}$ for $M = 1.8$ and as such will be under-sampled according the Nyquist theorem by a factor of two. But as one can see from the first paragraph of this section, the constraints do not allow enough freedom to correct this. To insist on the Nyquist limit at this magnification would mean decreasing the NA so much that the diffraction would limit the resolution to objects of 7 µm. But this is not as bad as it sounds. A feature 10 µm in size will still be imaged as three separate objects, each one about the size of one of our pixels. Since the features we want to examine are density fluctuations that rise and fall over length ℓ_ϕ, the wavelength of these fluctuations will be $2\ell_\phi \sim 20\,\mu\text{m}$ and will be adequately sampled.

4.3.1.1 Modification for In-Situ Imaging

For in-situ imaging, ToF is not a factor. The atoms remain in a trap narrow enough to confine them in its ground state and therefore, at a size of $2a_\perp \approx 2\,\mu\text{m}$ one can presume, inside of the DoF. Therefore everything can be optimized for resolution. By changing out the negative meniscus lens for a plano-concave one with a shorter focal length and enlarging the aperture one can obtain a high-resolution lens which works at the Nyquist limit by matching its magnification to its resolution via equation (4.11). Unfortunately this does require refocusing all of the optical elements[5] including the collimating group. This means that the front group is no longer perfectly collimating the light from the atoms and the distances the lenses take from the stop aperture are no longer arbitrary. Extra length between the telephoto element and the camera is also necessary, so while the changeover is not extraordinarily elaborate and requires two extra parts, it will require some imaging down-time and increase the overall length to 40 cm.

While there are lenses available with $D \approx 1$" with shorter focal lengths than the one I have chosen, they are so short as to introduce significant aberrations at the larger apertures needed for high resolution. The lens presented in table 4.3 and figures 4.7 and C.2 has $M = 6.4$ and $\text{NA}_{obj} = .1782$. Notice the ring of blue crosses just inside the black circle denoting the Airy disk in Figure 4.7 c). This suggests that the aperture cannot be enlarged any farther or the focal length of the plano-concave lens reduced without aberrations beginning to affect image quality. In fact Thorlabs list the effective pupil diameter (EPD) for their 2" lenses as 45.72 mm, just over this design's aperture of 44 mm. By Figure 4.2 (b) it can be seen the resolution will be diffraction limited rather than geometrically limited, as it should be at or above the Nyquist limit.

With $M = 6.4$ and $\text{NA}_{obj} = .1782$ the lens has a resolution of $\text{dy}_{min}(NA) = 2\,\mu\text{m}$ and an $M_{\max} = 6.398$. This means it is almost truly diffraction limited instead of sampling limited.

5. Compare table 4.2 and table 4.3 lines 5 7, 14, and 16.

In-situ objective

Surface	Radius	Thickness	Glass	Semi-Dia.
1 Object	∞	81.8	Vacuum	—
2 Window glass	∞	6	Silica	18.35
3 Window flange	∞	6.5		17.55
4 Bolt heads	∞	8		30
5	∞	3.965		25.4
6 250 mm meniscus lens	-224.67	5.53	N-BK7	25.4
7	-82.540	12.207		25.4
8 250 mm doublet	459.2	2	SF2	25.4
9	111.69	7.5	N-BK7	25.4
10	-137.09	10		25.4
11	∞	10		22
12 200 mm doublet	109.86	8.5	N-BK7	25.4
13	-93.11	2	SF2	25.4
14	-376.25	172.819		25.4
15 −75 mm singlet	-38.59	3.5	N-BK7	12.7
16	∞	152.958		12.7
17 (Camera opening)	∞	17.526		12.7
18 (Sensor)	∞	—		—

Table 4.3 − Surface specifications for high-magnification, high-resolution lens for in-situ imaging. It has a magnification of $M = 6.4$. Resolution is 2 µm and is sampled at the Nyquist limit. As this is not exactly an infinite system, no lengths are arbitrary. Thicknesses are measured on the optic axis and all units are mm.

Indeed, it makes an Airy disk with radius $r_A = 12.88\,\mu m$, very close to the ideal value of $2l_{px} = 12.9$. If one takes the MTF cutoff frequency as the minimum resolvable object instead of the Rayleigh criterion to indicate the smallest resolvable object one gets by equation (4.3) $\nu_c = 605\,mm^{-1}$ or a minimum object size of 1.65 µm.

There is a cost in terms of length, this lens being 10 cm longer than the one in table 4.2 and Figure 4.6. It also has a minuscule Dof of only 5.7 µm which would give a Tof of around 70 µs if it were used for that application. But solving equation (4.15) for $t = 0$ shows this is enough to image the in-situ cloud which would be only 3.3 µm wide.

4.4 Conclusions

The three objectives presented here will let us perform time-of-flight and in-situ imaging in the science chamber. This will let us collect most of the data described in Section 2.3, including the data from Bragg spectroscopy. A different objective will be needed for the light-sheet fluorescence imaging. That will aim at the lower window of the science chamber which has a diameter of 63 mm rather than 40 mm for the side windows. The light sheet is expected to be closer to this window (at most 40 mm away) than the atoms are to the side windows (just under 82 mm). This allows for an objective numerical aperture of up to 0.62 if the whole area of the window is used. This higher numerical aperture is useful, not only for resolution, but also because the atoms passing through the light-sheet will not scatter very many photons

(in [178] they expected to capture fifteen of the eight-hundred photons scattered by each atom).

The low-magnification objective presented here will have a field of view of $27 \times 20\,\mathrm{mm}^2$ in object space with no vignetting and allow an ample 10 to 15 ms ToF depending on the cloud temperature. For comparison, we currently us a ToF of 6 ms in the MOT chamber. Resolution is not critical but will be at worst (insisting on 50% contrast ratio, closed down aperture) 75 µm for an object with a size on the scale of millimeters. The adjustable iris allows some flexibility for trading off resolution for DoF.

The high-magnification lens here, may pose difficulties due its short depth of field but it offers a certain amount of flexibility. If the high-focus objective turns out to be too shallow one can refocus the rear group[6] to get $M = 1.6$ while aberrations expand the beam just to the boundary of the Airy disk on fields 1-3, but this only yields a DoF of 38.4 µm and a ToF of 5.8 ms. In principle, the SM1V10 can be extended 2 ¢m and the rear lens removed altogether giving $M = 1.4$, DoF = 43.9 µm, and ToF = 6.6 ms. But these minor extensions of the ToF are paid for by a loss of resolution ($\mathrm{dy}_{min}(M) = 3.6\,\mu\mathrm{m} \to 4\,\mu\mathrm{m}$ or 4.6 µm, respectively) so one must hope that optimism about the assumptions behind the geometrically limited resolution (4.6) is justified. Fortunately all of these adjustments are minor and do not require new lenses.

Only the low-magnification objective has been constructed but it has not been installed or tested since they were intended for the science chamber which has only just been brought online. However, by the calculations presented here, these lenses should be able to gather the data we need.

6. In table 4.2 lines $14 : 130.305\,\mathrm{mm} \to 159.188\,\mathrm{mm}$ and $16 : 61.894\,\mathrm{mm} \to 19.803\,\mathrm{mm}$

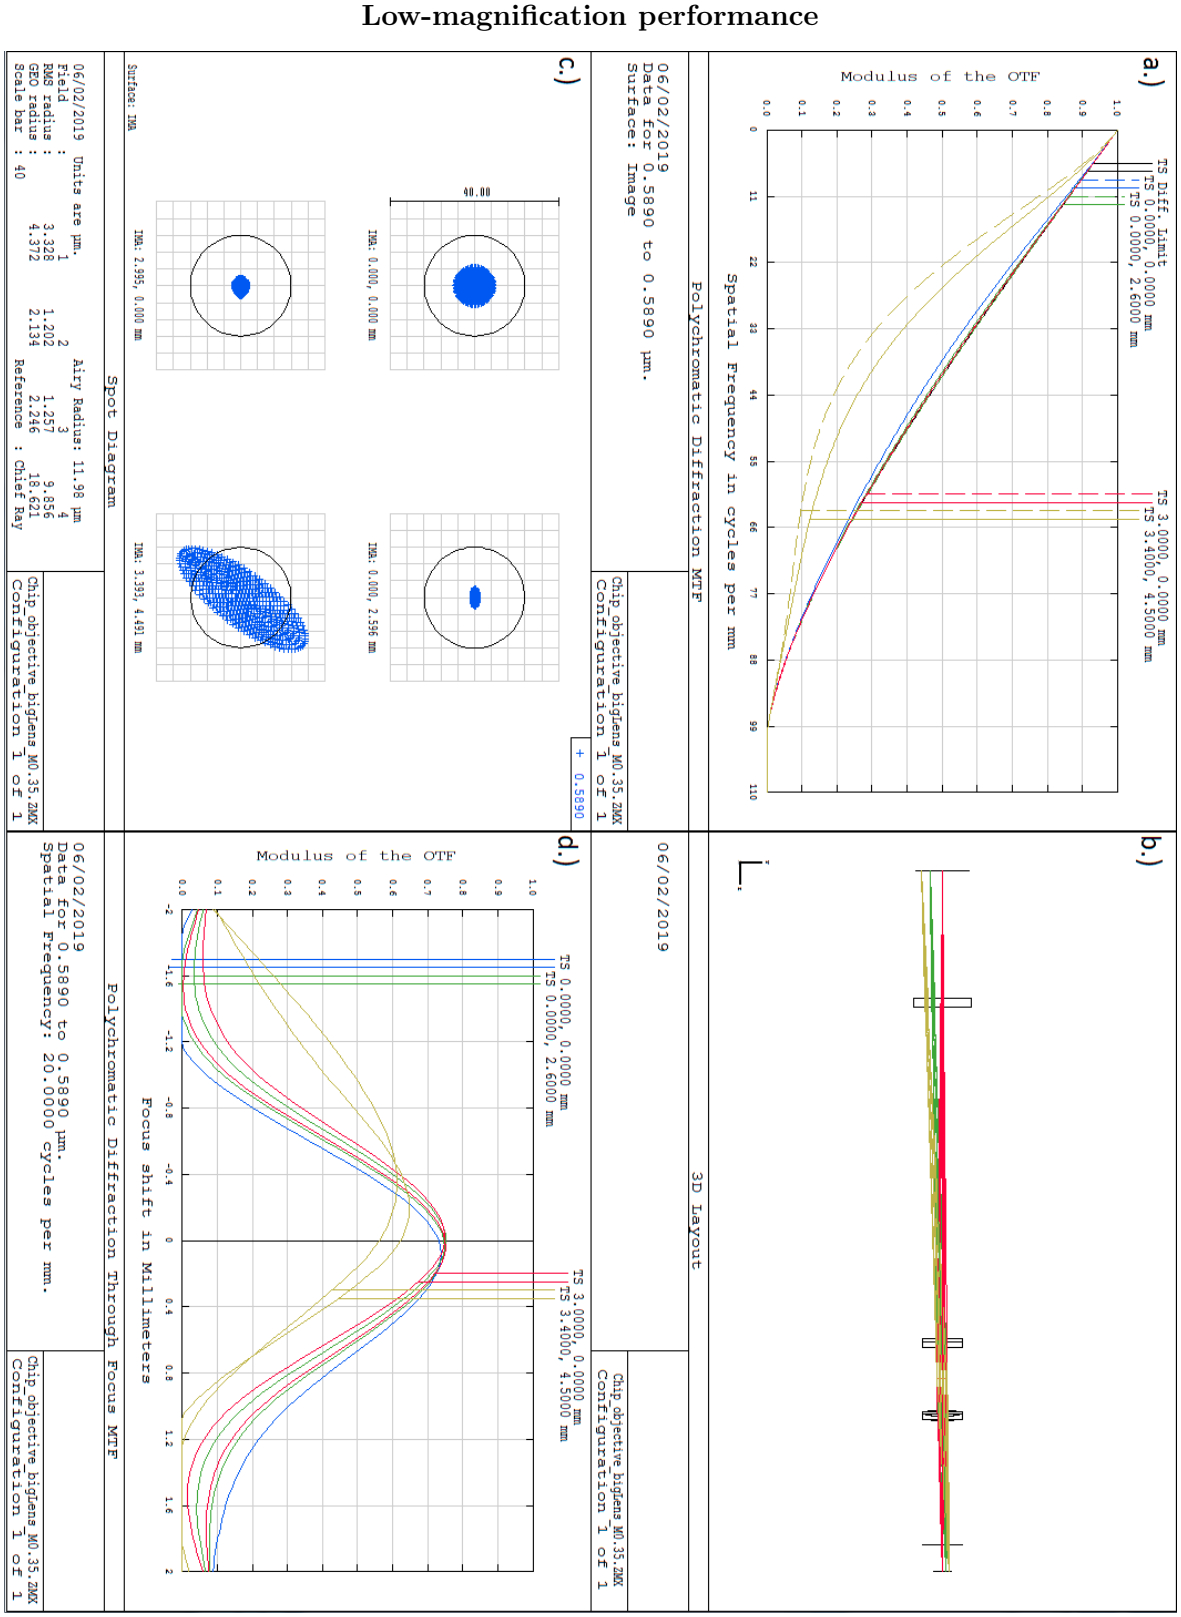

Figure 4.5 — Design and performance data for the low magnification objective at aperture radius $r = 3\,\mathrm{mm}$. a.) FFT MTF, b.) Layout of lens. c.) Spot diagram. d.) Through-focusFFT MTF at $20\,\mathrm{mm}^{-1}$. The fields are at the following (x, y) coordinates in image space, taking the optic axis in the center as the origin: field 1: on-axis. 2: $(0, 2.6\,\mathrm{mm})$, 3/4 of the way up to the short edge of the sensor. 3: $(3\,\mathrm{mm}, 0)$, 2/3 of the way to the long edge. 4: $(3.4\,\mathrm{mm}, 4.5\,\mathrm{mm})$ at the sensor's corner.

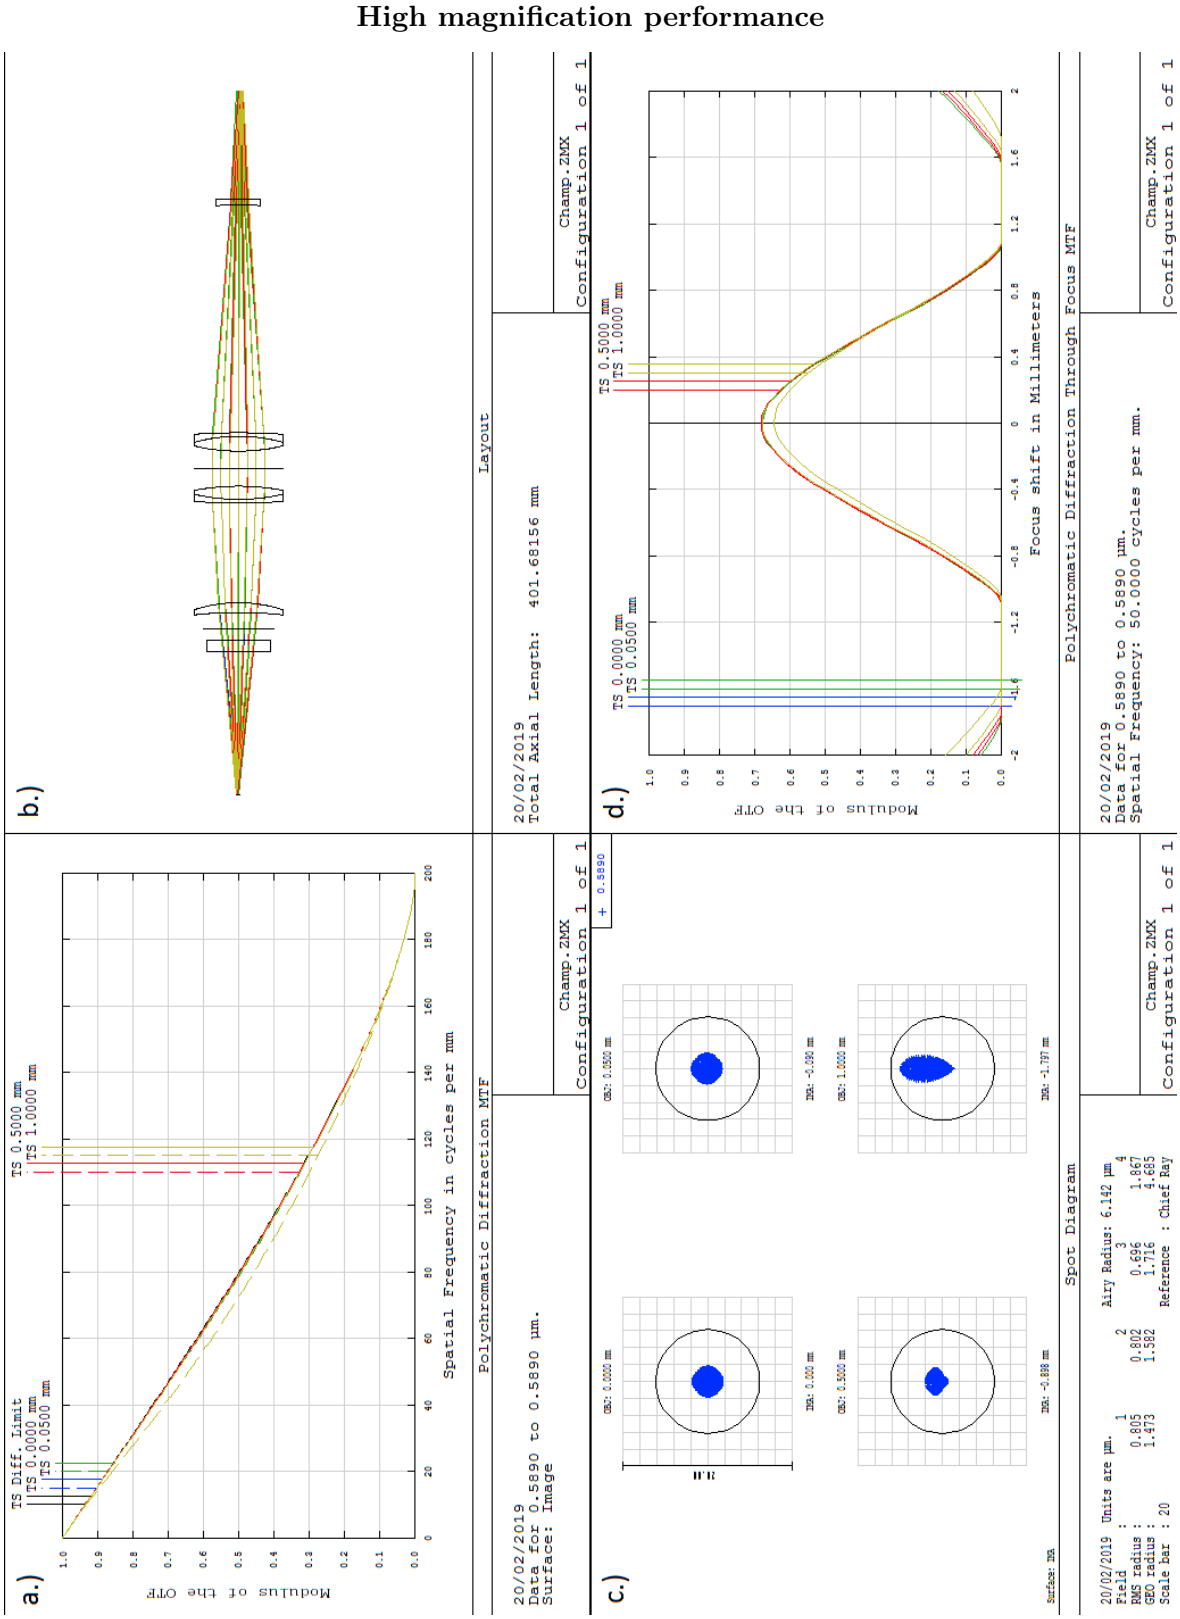

Figure 4.6 — Design and performance of the "high" magnification system with $M = 1.8$. a.) FFT MTF, b.)ray diagram, c.) spot diagram, d.) through-focus FFT MTF at $50\,\mathrm{mm}^{-1}$ or a 10 µm object size. Field 1 on-axis. Field 2 = 50 µm off axis in object space $\implies$ an object 100 µm across. Field 3 = 500 µm. Field 4 = 1 mm. The through-focus MTF, d.), is misleading because this is the difractive DoF which is much greater than the geometric one.

In-situ objective performance

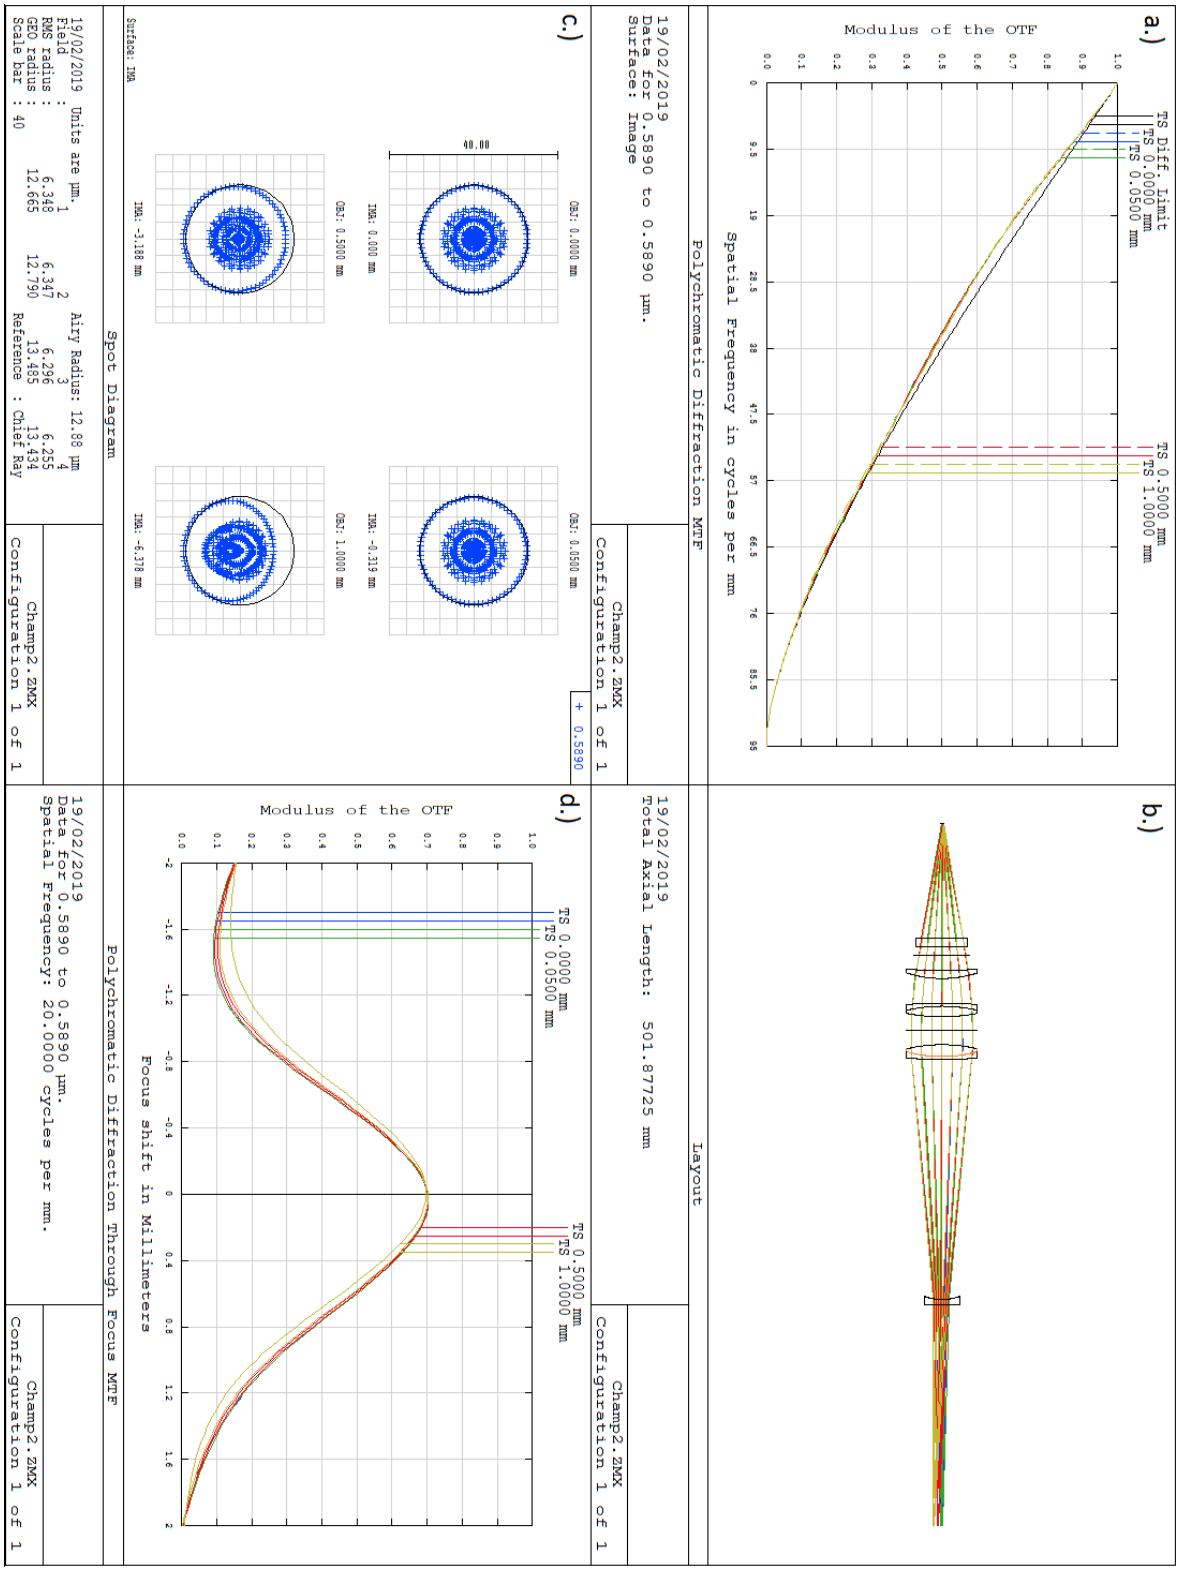

Figure 4.7 — Design and performance of the in-situ imaging system. a.) FFT MTF, b.) ray diagram, c.) spot diagram, d.) through-focus FFT MTF at 50 mm^{-1}. Field 1 = 0 µm, field 2 = 50 µm off axis in object space $\implies$ an object 100 µm across. field 3 = 500 µm field 4 = 1 mm

Chapter 5

Magnetic transport optimization

When I arrived, the experiment up until the MOT chamber had been entirely assembled [1]. The transport up to where it would meet the science chamber was mounted and the system up to the valve at the corner of the transport (marked in figure 5.4) was under vacuum. The rest of the system was yet to be installed. While dealing with various issues related to the chip and its mount, we undertook to optimize the transport of the cloud in the magnetic trap by taking it from the MOT chamber, up to the corner of the transport, and back for imaging. Much work has been done on transport of cold atomic clouds: via optical tweezers [230–232], an anti-Helmholtz coil pair on a mechanical translation stage [233, 234], one-dimensional optical lattices [235], chains of various magnetic traps generated by chips [236, 237], chains of Ioffe–Pritchard traps [238], and conveyor belts of anti-Helmholtz coils such as we use [2, 239, 240]. These techniques are necessary whenever the cloud is formed in one part of the experiment and then some other process, such as condensation or probing, is performed in another. To take our experiment as an example, it allows us to load our MOT while performing experiments in the science chamber and so increase our repetition rate and protect atoms in the science chamber from stray resonant light. This is especially useful when studying fluctuations, as we are, since many repetitions are needed to gather meaningful statistics.

Optimization work has been done on transporting atoms in harmonic [241–243] traps and harmonic traps with higher-order power law terms [244, 245]. We cannot rely on the techniques used in those studies because they require the potential to have a well-defined derivative everywhere. We are working with a trap with a linear potential so the derivative is not well-defined at the minimum. Instead we simulate moving a trap via various time profiles and calculate the trajectories of classical particles under its influence. We then experimentally test these profiles, using the results from one method to inform what to test with the other.

The state-of-the-art transport profile is one where the trap center is moved at a constant velocity of less than $15\,\mathrm{cm}\cdot\mathrm{s}^{-1}$, to avoid heating, resulting in a retention of between 60 and 80% of the initially trapped atoms [2,239,240]. The loss mechanism is presumed to be collisions with the background gas. This transport involves heating the cloud by around 15% [2,239] or 50 to 70% [240] (who begins at a temperature an order of magnitude lower). Other transport profiles for quadrupole trap chains have been tried [240], but did not provide an improvement in average speed for comparable loss and heating rates. By careful choice of time profile we were were able to achieve an average speed of over $50\,\mathrm{cm}\cdot\mathrm{s}^{-1}$ while keeping over 70% of our atoms and heating the cloud $\sim 20\%$. Via our simulations we were also able to say which loss

mechanism was more prevalent [246].

In the first section I will give the relevant theory concerning quadrupole traps and using chains of them for transport as well as define the profiles I will discuss. Then I will go into a few technical details: the parameters of our particular trap, the limitation on our holding time in the magnetic trap in the MOT chamber, the collision rate in the trap which gives the thermalization time, and the various methods we use to determine the temperature and currents. Then I will discuss the experiment, the simulations and compare their results. Finally I will draw my conclusions.

5.1 Theory and background

As discussed in 3.1.1, neutral atoms will still have a magnetic moment depending on their spin state. This means that, by certain arrangement of fields and spin states they can be subject to magnetic forces which can confine them. First I will talk about these circumstances and the particular field arrangement we use to confine them, the quadrupole field. Then, given the potential felt by atoms in such a field, I will talk about their distribution when they have reached thermal equilibrium. This trap is subject to certain kinds of losses, some related to how the trap is moved, and so relevant to our optimization, some not. I will talk about which is which and go into discuss the relevant ones.

Then I will discuss how this trap is to be moved: by a chain of coils similar to the ones that make the stationary trap. I start from the example of two coil pairs to introduce the effects this strategy has on the geometry of the trap as a way of introducing the three-coil method we employ. This is the maximum number of coils which need simultaneous current to control the trap parameters. This sets up the details of how and when these three currents are switched between sets of three coil pairs.

Lastly I will define the three current profiles whose results paint the clearest picture of our investigation. One is the standard, constant velocity method. One provides a foil to it, solving its particular problems and introducing others. A third provides a compromise. For other profiles we tried see Appendix D.

5.1.1 A single quadrupole trap

As explained in § 2.2.4, atoms in a weak magnetic field, $B(\boldsymbol{r})$, will feel a potential of

$$V(\boldsymbol{r}) = g_F \mu_B m_F B(\boldsymbol{r}) \tag{5.1}$$

and be either low-field-seeking or high-field seeking depending on if the product of the Landé g-factor, g_F and the spin projection, m_F, is positive or negative, respectively.

The field density, and therefore the potential in (2.7), reaches a maximum in between the actual loops, where the distance to the current is minimal. At this distance from the center in the horizontal plane, the trap breaks down. Since the horizontal gradient is half that of the vertical one, this is more important for determining how much energy an atoms needs to escape the trap. This energy defines the trap depth and atoms lost by reaching the region where the trap breaks down are said be lost via *finite-depth losses* which I discuss on page 104.

5.1.1.1 Cloud Properties

From the potential energy (2.7) we can derive the properties of a cloud in the trap at thermal equilibrium. At thermal equilibrium (and neglecting any interaction with the background gas), the atoms will be Boltzmann distributed, where the probability for being in a state indexed by i, will be

$$p_i \propto e^{\frac{-E_i}{k_B T}} \tag{5.2}$$

where E_i is the energy of the state, k_B is the Boltzmann constant, and T is the temperature of the system. Since the position-dependent energy is the potential energy, the number density of atom at position $\boldsymbol{r}$ will be

$$n(\boldsymbol{r}) = C \exp\left(-\frac{V(\boldsymbol{r})}{k_B T}\right) \tag{5.3}$$

were C must be the total atom number N divided by the effective volume of the cloud, V_{ol}, since integrating n over the volume must yield N.

To simplify the expression for the volume, among other things, we can define a length scale for the density,

$$R = \frac{k_B T}{\kappa}. \tag{5.4}$$

At $T = 250\,\mu\text{K}$, an unusually high but not unknown temperature for our cloud, $R = 1.15\,\text{mm}$.

Then to find V_{ol}, we start from equation (5.3) and write $V_{ol} = N/C = \int n(\boldsymbol{r})\,\mathrm{d}\boldsymbol{r}$. A coordinate change $x' = x$, $y' = y$ and $z' = 2z$, implying a $\boldsymbol{r}' = (x', y', z')$ before changing to spherical coordinates lets us perform the integration

$$V_{ol} = 2\pi \int r'^2 \exp\left(-\frac{\kappa r'}{k_B T}\right)\mathrm{d}r' = 4\pi R^3 \tag{5.5}$$

$$\implies C = \frac{N}{4\pi R^3}\,. \tag{5.6}$$

5.1.1.2 Losses

Atoms in a magnetic trap are subject to a range of losses: there are spin-exchanging collisions which can throw atoms into untrapped spin-states [169]. There are collisions with the background gas that can eject the atoms from the trap, the background gas being at a much higher temperature than the gas in the trap. There are N-body recombination losses (see § 2.3.2) where three or more atoms interact to give one of the atoms enough kinetic energy to escape the trap. There are finite-depth losses where the atoms find themselves in a region where potential no longer confines them. And there are Majorana spin-flip losses in region around the magnetic zero. These last two are the only ones affected by the trajectory of the trap center, I will go into them in some detail here. But first I need to justify why the other loss sources do not depend on the trap trajectory.

Spin exchange collisions are interactions where the total angular momentum vector, $\boldsymbol{F} = \boldsymbol{F}_a + \boldsymbol{F}_b$, of the interacting atoms, indexed a and b, is preserved, but the angular momentum of the individual atoms is not. This can lead to changes in the hyperfine levels and tends to align the polarization of the atoms [169]. Two atoms initially in $|F = 2, m_F = 1\rangle_{a,b}$ can scatter into $|m_F = 0\rangle_a + |m_F = 2\rangle_b$, and one of the atoms will be lost. We avoid this by having only $|F = 1, m_F = -1\rangle$ atoms in the trap. In that case, two atoms colliding would have total

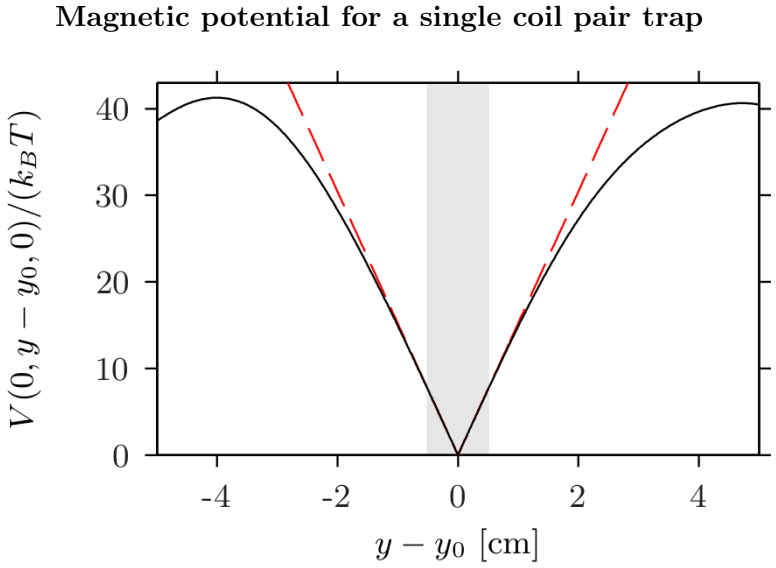

Figure 5.1 — Potential generated by a single transport coil (e.g. C_5 in the left of Figure 5.4) along one axis. Solid line is the analytic formula given by [170], dashed line is the linear approximation fit over the shaded region. Notice how the potential falls away from linear near the inner radius of the coil (1.33 cm) and begins not to trap at all near the outer radius (3.97 cm).

spin projection $m_F = -2$ which cannot be the spin projection of a single atom in $|F = 1\rangle$ so they must stay in the $|F = 1, m_F = -1\rangle$ trapped state.

Collisions with the background gas are unavoidable, absent a totally perfect vacuum. The rate of loss depends on the relative velocity between the trapped atoms and the background gas, so transporting the atoms, could, in principle, affect this loss rate. However, since the background gas is at room temperature of 293 K, the average air molecule has a velocity of $\sqrt{k_B T/m_a}$ where m_a is the mass of the air molecule. This means that even the relatively slow N_2 molecules have an average speed of close to $420\,\mathrm{m}\cdot\mathrm{s}^{-1}$, far faster than we hope to transport our atoms. So we neglect the effect the transport profile will have on the loss rate due to background gas collisions.

The only recombination losses we face are 3-body recombination losses, since there is no light to induce photoassociation. Two atoms colliding can only exchange momenta and need a third to impart their energy to if they are going to form an untrapped molecule and higher-order losses are even less likely than thee-body ones. But this means that three atoms need to meet at once, which means multiplying by another factor of the density (among other things, see § 2.3.2 and references therein). In the magnetic trap the density is so low that collision rate is already on the order of once per second (see Figure 5.6) so three-body losses are not expected in an experiment lasting $\lesssim 10$ s.

Since these losses are not significantly affected by how the cloud is transported but are still present, they need to be discounted by normalizing the atom number after transport by the atom number remaining in a trap that was not transported but rather held in the MOT chamber for the transport time. See § 5.3.1 for more details.

This leaves finite-depth losses and Majorana spin flip losses which cannot be normalized out and do indeed depend on the transport profile.

Finite depth losses Since the quadrupole magnetic field is made by a pair of coils of finite size, the approximate linear potential described in (5.11) is only valid up to a certain distance from the trap center, as illustrated in Figure 5.1. In between the two coils, the magnetic flux lines will be densest and therefore the magnetic field strongest; past that and they begin to diverge (see Figure 2.4). In the region outside the coils, therefore, the radial gradient becomes negative and the potential is no longer trapping. So, if an atom finds itself far enough from the center it will be lost. This constitutes a finite-depth loss; the depth of the trap giving the maximum kinetic energy an atom can have and still be held in the trap. For an idealized pair of coils, having only one loop, this will be given by the potential energy in the horizontal plane passing through the magnetic zero, at the radius of the loops. Since the potential remains roughly linear up this point, falling off slightly before [170, Fig. 3(b)], this energy, V_d, can be estimated by

$$V_d \lesssim \kappa\rho \tag{5.7}$$

where ρ is the radius of the loops and $\kappa = g_F m_F \mu_B b'$ as above. This can also be given as a temperature, $T_d = V_d/k_B$. For our coils with multiple loops, one can estimate V_d by the linear approximation at the average radius over all the loops (see Figure 5.1).

Once the trap has reached thermal equilibrium, the cloud will be at a temperature $T < T_d$.[1] With atoms that would have raised the temperature over T_d having been lost. However, high acceleration can move the trap far enough, fast enough to put atoms outside of where the trap can contain them. This very much depends on the trap trajectory and is neatly dealt with by considering an effective potential in a frame co-moving with the center of the trap, $\tilde{V}(\boldsymbol{r})$. Since the gradient added by the moving frame is just $m\boldsymbol{a}$ we may write

$$\boldsymbol{\nabla}\tilde{V}(\boldsymbol{r}) = \boldsymbol{\nabla}V(\boldsymbol{r}) + m\boldsymbol{a} \tag{5.8}$$

where $\boldsymbol{a}$ is the acceleration of the trap center. This leads to a trap depth in the moving frame $\tilde{V}_d = \min_{\boldsymbol{r}} \tilde{V}(\boldsymbol{r}, a)$. Then if the magnitude of the acceleration, $\|\boldsymbol{a}\|$, ever goes over $\boldsymbol{\nabla}V/m$, $\tilde{V}_d$ will fall bellow zero and the atoms will not be confined at all by the trap on the side it is accelerating away from. Any atoms that are too far away from the trap center to be trapped when the acceleration diminishes will be lost. How many atoms are lost then depends on how large the acceleration is, and for how long it persists for. This last caveat is why the usual constant velocity profile does not lose all the atoms even though it has an infinite acceleration spike.

Majorana Losses The other kind of losses that can be influenced by the trap trajectory are Majorana or spin-flip losses, introduced in § 2.2.4. Because these happen when an atom sees a weak field changing quickly, they are affected by the trajectory of the trap. Equation (2.9) lets us calculate a region where Majorana losses are expected to occur. By taking $p_{\mathrm{Maj}} \to 1$ and solving for $\boldsymbol{r'}$ we can get the radius of an ellipsoidal region where losses are expected. This is useful for calculating loss rates both a priori to experiment [1, §4.1.3.2] and in simulations.

The trap trajectory will determine where in the trap the atoms are located over the course of the transport. This is the primary way the trajectory influences the rate of Majorana losses: by placing the atoms near or far from the region mentioned above where Majorana losses are likely. In principle, they could also the atoms quickly through this region, further increasing losses. This did not seem like a major cause of losses for our profiles, with the exception of the "constant acceleration with a pause" profile mentioned in Appendix D, which we suspect

1. In practice, $T \ll T_d$, cooling having occurred in the MOT.

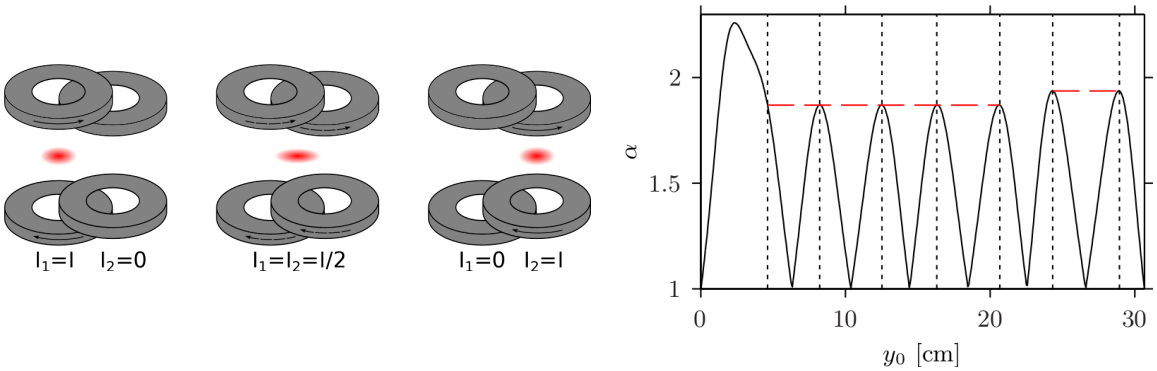

Figure 5.2 — A transport of only two coils must stretch and squeeze the trapping potential as they pass the cloud between them. Left — Cloud as transported by two coils. Adapted from [1]. Right — Aspect ratio, α of the trap produced by our transport using only two coils. Large peak at the beginning is because the MOT chamber coil has a larger radius.

suffered from this mechanism. The loss rates where so severe that we did not make a thorough investigation with simulations.

5.1.2 Quadrupole trap chains

Chains or conveyor belts of anti-Helmholtz coils are a standard way of moving clouds of cold atoms [2, 239, 240]. Half-overlapped coils produce quadrupole fields which sum to a trap that can be passed down the chain by leapfrogging which coils current is supplied to. For details about the construction of our specific transport chain, including characterization of its field, see [1, Ch. 4]; this chapter deals with optimizing the transport and I will confine myself to material relevant to that.

Consider first a pair of two half-overlapping[2] anti-Helmholtz coil pairs with their axes arranged on the y-axis as in Figure 5.2. All of the reasoning from the last section about how this must create a magnetic zero in the region between the coils still holds but now the symmetry around the axis is broken when current is flowing through both of them. When all of the current is in one coil, the potential is symmetric as in equation (2.7) and the minimum is at the center of the coil with all the current. When the current is equally distributed between the two pairs the net current where the coils overlap is effectively zero, the two counter-flowing currents canceling each other out. This puts the point where $B = 0$ between the axes of the two coils; making the two circular coils behave like one elongated coil, producing a lower magnetic gradient in the y direction and elongating that axis of the trapped cloud. To deal with this asymmetry in the trapping potential let the ratio of the horizontal gradients be

$$\alpha \equiv \frac{b'_x}{b'_y} \tag{5.9}$$

where b'_x and b'_y are the gradients along x and y respectively. This makes the magnetic field

2. Meaning the axis of one coil coincides with the outer radius of the other.

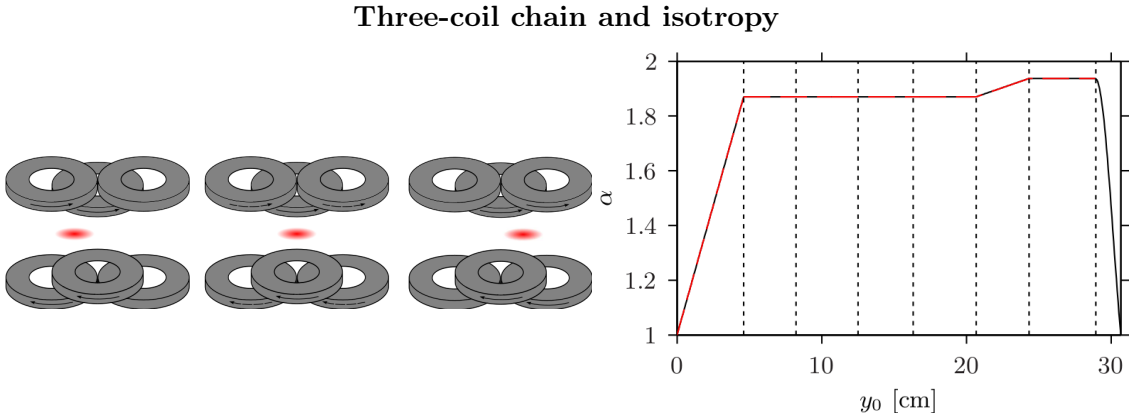

Figure 5.3 — Similar to Figure 5.2 but for three coils. Left figure again adapted from [1]. Notice that while α can be held constant, it cannot be decreased, with the exception of the first peak while leaving the MOT chamber. This is thanks to the *push coil* near the MOT (see 5.4).

read

$$\boldsymbol{B}(x,y,z) \simeq 2b' \begin{pmatrix} -\dfrac{\alpha}{1+\alpha}x \\ -\dfrac{1}{1+\alpha}(y-y_0) \\ z \end{pmatrix}, \tag{5.10}$$

where b' is now half the gradient along z, which we try to keep constant, and α handles the asymmetry. Note that in this scheme, assuming the currents are all of the same polarity, the trap is only ever stretched along y and simultaneously compressed along x so $\alpha \geq 1$. If negative current could be supplied to the coils, the gradients along y could be set freely and $\alpha = 1$ could be maintained. By equation (2.5) the field (5.10) results in a potential of

$$V(x,y,z) \simeq 2\kappa\sqrt{\frac{\alpha^2 x^2}{(1+\alpha)^2} + \frac{(y-y_0)^2}{(1+\alpha)^2} + z^2}\,, \tag{5.11}$$

where $\kappa = g_F \mu_B m_F b'$. Varying the balance of current between the two coils can shift the magnetic zero back and forth between them.

This makes it clear that, with a chain of enough overlapping pairs of coils, a cloud of magnetically trapped atoms can be moved arbitrarily far. That is to say that y_0 can be controlled. By changing the total amount of current, the magnitude of the gradient, b' can be scaled. It is also clear that the horizontal anisotropy of the trap will change, but with only two coil pairs, carrying current in the same direction, there is no way to *control* how α evolves interdependently from y_0. It must fall to one every time the center of the trap reaches one of the coil axes and rise above it as y_0 moves between them, reaching a maximum when it is equidistant from the axes as seen in Figure 5.2. This constant squeezing and releasing of the cloud will lead to heating. This can be remedied by adding a third coil whose axis is at the outer radius of one of the pairs, Figure 5.3. This puts the atoms in a situation analogous to being constantly held between two elongated loops, never being held by a single, circular coil pair. Thus the trap is elongated but, if the currents are chosen correctly, it is consistently elongated and the aforementioned heating is avoided.

Determining the currents

From the last section it is clear that the moving trap can be characterized by three parameters: the position of the magnetic zero, y_0, the magnetic gradient, b', and the horizontal anisotropy, α. Since the top and bottom coils in an anti-Helmholtz pair are linked in series and carry the same current, there is no control over the off-axis position in the x, z-plane. These three parameters can be controlled independently by the three currents if these three currents can be arbitrarily chosen. However, to get a steep enough gradient to trap atoms at the temperature reached by our MOT (see Chapter 2) currents up to 100 A are needed (see 5.4). High-current, bidirectional power supplies that would allow us to freely control all three parameters of our trap are not impossible to fabricate [247], but they tend to have been developed for special applications (see [248] and references therein) So we use more commonly available unidirectional supplies (see § 2.2.6 for details) capable of supplying up to 60 V and 100 A. But this means that we cannot arbitrarily control all three of our parameters and must make some compromises. For further details and a simulation of the magnetic field see [1, §4.2.3.3].

At some point in the process, the current must be switched from the pair of coils on the trailing side to the new member of the active triplet on the leading side. At that point you only have two currents and at that position the other two parameters, α and b', cannot both be freely set. Fixing b' has the advantage that it yields current profiles where the current is always positive. There is a further constraint that these switches should always occur when the current in the hindmost coil has fallen to zero. Otherwise there is a risk that the sudden change in current will lead to eddy currents (see [1, §4.4.2.1] for more about our efforts to minimize eddy currents) which will distort the trap and could heat the cloud. So α is the last remaining parameter which can absorb the incompleteness of our control.

Because our power supplies only supply positive currents, $\alpha = 1$ cannot be maintained as the trap moves and must rise above unity as explained before. Imagine a situation where the trap is being generated by three coils simultaneously. If all three currents are positive, all three coils will tend to drag the magnetic field minimum to their respective centers; the only control we have is over the amount of current and therefore the degree to which each coil pair wins out. This will stretch the trap along the axis of the chain and raise the value of α. Since we are trying to avoid having the cloud expand and contract along the axis as we move it, and turning on another coil with positive current can only raise α, our strategy was to keep α at the peak values shown in Figure 5.2 giving the profile shown in Figure 5.3. The exception being in the transfer between the first pairs where we use the *push coil* shown in Figure 5.4. Since positive current through that coil tends to lower α, the switch point comes when α matches the maximum value of the *next* pair. The reason α climbs so high during that first transition without a push coil is that the coils that make the magnetic trap in the MOT chamber are larger than the coils in the transport chain, having an inner radius of 36.2 mm versus 13.4 mm [1, Tbl. 4.2]. The anisotropy must fall to one at the beginning and end of the transport since the trap is being generated by only one coil there. At the end of the transport going into the final coil, there are only two coils available and α rises to a value determined by the size (and therefore spacing) of our coil pairs. We could have added extra coils at the corner, to always have three coils available and gain better control over α as in [240], but geometric constraints in our experiment prevented us.

For calculating the current profile which meets these conditions, we use an analytic expression for the field generated by a pair of single loops of a given diameter and separation, carrying a given current [170]. From this we construct an expression which accounts for the winding of the coils, given an inner and outer radius, and use that to calculate the trap pa-

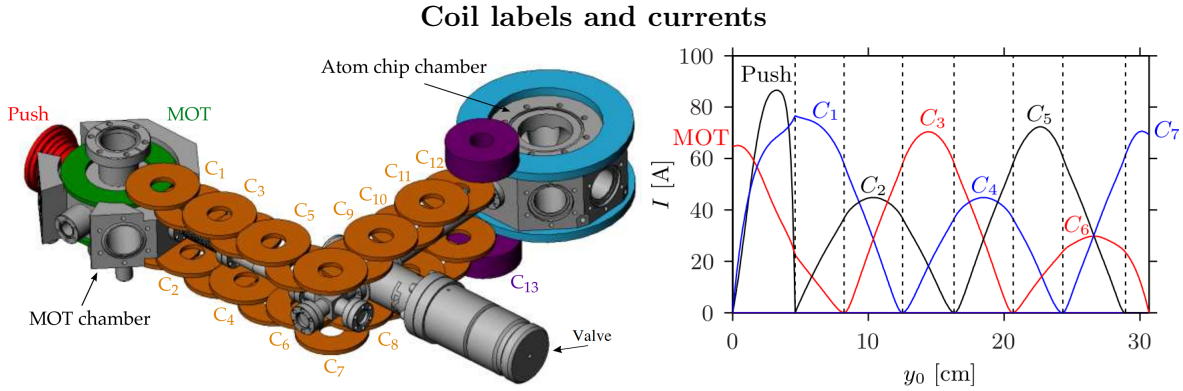

Figure 5.4 — Left — Labels and locations of the coils in the magnetic transport. Right – A current profile for a one-way transport. Switches between coil triplets marked with vertical dashed lines, each power supply is plotted in its own color. Notice switches occur when the current falls to 0 A. Some currents are larger because those coils are farther away vertically so they need more current to keep b' constant.

rameters yielded by different current profiles. This expression holds far from the trap center (extending even beyond the coils themselves) where (5.11) does not. They were used to generate the solid black line in 5.1. They do not take into account the helicity of our coils, but because they are only two layers packed flat against each other that is not expected to be significant.

In order to make sure our approximate formula for the trap potential given in equation (5.11) would always be applicable to our atoms, it was used as a fitting model for the analytical formula. The potential was calculated at 101 points along each of the axes, each set of 101 points spanning ±5 mm from the trap minimum at $\boldsymbol{r}_0$. A least-square algorithm was then used to minimize the difference between the potential given by the analytic formula and the one given by the approximate formula (5.11) at these points. The range covered by these points in the area highlighted in gray in Figure 5.1 and the difference at these on-axis points was less than than 0.1%. Since the radius of a cloud in the trap (5.4) at a typical temperature of 200 µK is $R \approx 1$ mm, the approximation (5.11) should hold good for our experiment.

5.1.3 Time profile definitions

Here I will lay out the definitions and notation used for the various time profiles of the magnetic trap position and define the three profiles that I will treat in detail in this chapter: the constant-velocity profile, the constant-acceleration profile and the error-function profile.

The position of the trap center is at $\boldsymbol{r} = (x_0, y_0, z_0)$ and the axes are oriented such that the trap center travels along the y-axis. It is taken from the MOT chamber to the corner of the transport, coil C_7 in 5.4, and then back to the MOT chamber where it is imaged. This constitutes a *run*. When the cloud is imaged at the end of the run, it is integrated along the axis it traveled and imaged in the x, z plane, with z being vertical. Let the distance from the center of the MOT chamber at the beginning of the transport to the center of the C_7 coil be L_1. The trap center will cover this distance in a one-way travel time of Δt.

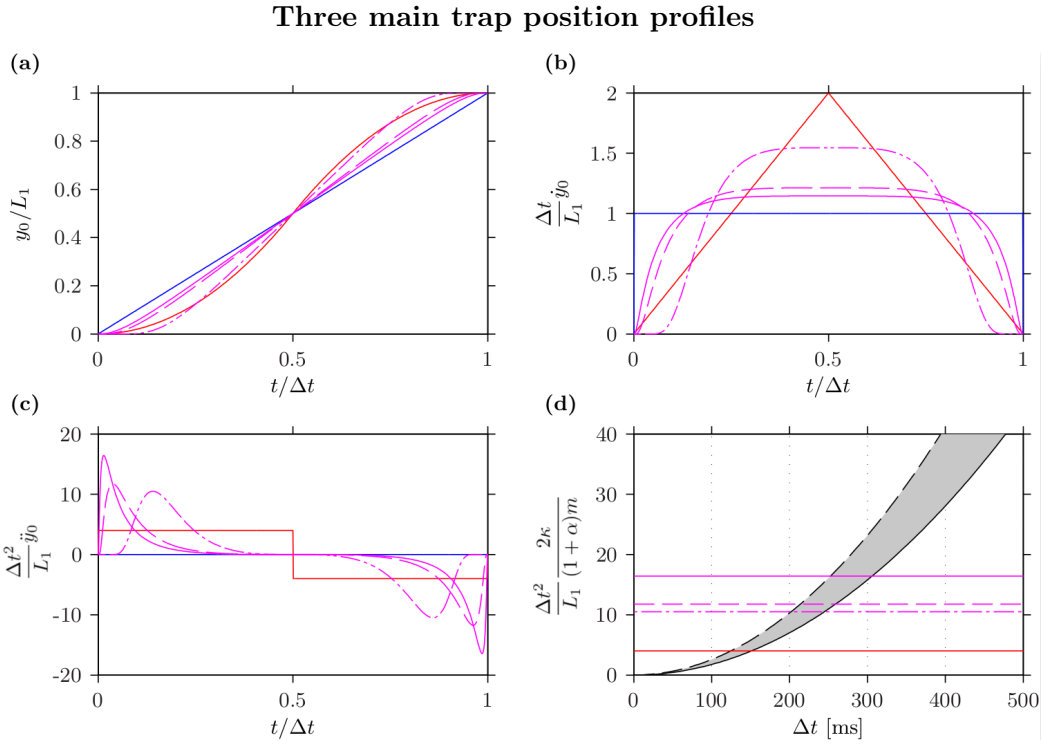

Figure 5.5 — Position (a), velocity (b), and acceleration (c) for the three time profiles. Constant velocity is shown in blue, constant acceleration in red and the three magenta lines correspond to the error function at three different values of the shape parameter, $\delta = 0.02$ (solid line), $\delta = 0.1$ (dashed), $\delta = 0.5$ (dash-dotted). The magnitude of the acceleration at which V ceases to trap for $b' = 65\,\mathrm{G}\cdot\mathrm{cm}^{-1}$ and $\alpha = 1$ (black dashed line) and our maximum value, $\alpha = 1.937$ (black solid line) is plotted in (d). The horizontal lines plot the maximum acceleration for the various profiles.

This means the trajectory of the basic constant-velocity profile reads

$$y_0(t) = \begin{cases} 0 & \text{for } t \leq 0 \\ L_1 \frac{t}{\Delta t} & \text{for } t \in]0, \Delta t[\\ L_1 & \text{for } t \geq \Delta t \end{cases} . \tag{5.12}$$

During the transport, when $t \in]0, \Delta t[$, the velocity is a constant $v = L_1/\Delta t$. However, at the beginning and end of the transport, when $t = 0, \Delta t$, the velocity suddenly changes from or to zero, giving a spike in acceleration as close to infinite in magnitude as is practically possible. To estimate how high this spike is, remember that a coil can be modeled as an LR series circuit of some finite inductance, L_c, and resistance, R_c having a time constant $\tau_c = L_c/R_c$ [216]. This means the current, and therefore the position of the trap center, cannot change more quickly than approximately $\delta t = 3\tau_c$. Taking the values from [1, Tbl. 4.1] for one of the typical transport coil pairs gives $\delta t = 1.5\,\mu\mathrm{s}$. For a profile of constant velocity v, the infinitesimal change in the velocity $\mathrm{d}v = v\,\mathrm{d}t$. The first instantaneous acceleration is then $\mathrm{d}v \to v/\delta t$ which is both finite and dependent on v, helping explain why fewer atoms are lost when the trap is moved more slowly.

To avoid dumping that atoms out this way, we also tested a profile with constant acceler-

ation, written

$$y_0(t) = \begin{cases} 0 & \text{for } t \leq 0 \\ 2L_1 \left(\frac{t}{\Delta t}\right)^2 & \text{for } t \in \left]0, \frac{\Delta t}{2}\right] \\ L_1 \left[1 - 2\left(1 - \frac{t}{\Delta t}\right)^2\right] & \text{for } t \in \left]\frac{\Delta t}{2}, \Delta t\right[\\ L_1 & \text{for } t \geq \Delta t \end{cases}. \tag{5.13}$$

This has an acceleration of constant, finite magnitude equal to $4L_1/\Delta t^2$ that switches from positive to negative halfway through at $t = \Delta t/2$ bringing the trap to a stop at $y_0(\Delta t) = L_1$. This has a large value of the jerk at the point where the acceleration changes direction.

As I will discuss in Sections 5.3.2, this is not the drawback in this particular profile, but to see how a profile smooth in all its derivatives performs, we tested one based on the error function,[3] written

$$y_0(t) = \begin{cases} 0 & \text{for } t \leq 0 \\ \frac{L_1}{2}\left[1 + \operatorname{erf}\left\{-\gamma\left(\frac{t}{\Delta t}\right)^{-\delta} + \left(1 - \frac{t}{\Delta t}\right)^{-\delta}\right\}\right] & \text{for } t \in]0, \Delta t[\\ L_1 & \text{for } t \geq \Delta t \end{cases} \tag{5.14}$$

where $\delta > 0$ is the tunable shape parameter. The coefficient $\gamma = 2^{-3/2-\delta}\sqrt{(\delta+1)(\delta+2)}/\delta$ is a function of δ that ensures that the jerk of the trap center $\dddot{y}_0(t) \to 0$ halfway through the transport at $t = \Delta t/2$. Since this is also where the acceleration changes sign, and this profile is smooth in all its derivatives, (5.14) tends to have a low acceleration over most of its length (see figure 5.5).

When the trap is expected to be anti-trapping due to its acceleration is shown in Figure 5.5 (d). The gray shaded area corresponds to the range of the absolute value of the slope potential in the lab frame, V, over the course of a run wherein α varies between $\alpha = 1$ (the higher value marked by a dashed line) and $\alpha = 1.937$ (the lower value marked by a solid line). As long as this range is above the maximum absolute value of the acceleration, $\tilde{V}_d > 0$ and the potential will be trapping.

5.2 Thermodynamic properties of the cloud

Here I go over a few details that will be useful for understanding our experimental results: the techniques we use for determining the temperature, and the collision rate in the trap. For details about the experimental control and imaging setup see Chapter 2 and [1].

5.2.1 Temperature measurement

The temperature of the cloud is measured when it is released from the stationary, and the therefore symmetrical, trap in the MOT chamber. We have three methods for making this measurement: fitting the growth of the variance of several ToF pictures; fitting the doubly-integrated, one-dimensional profile from a single image; and fitting with a two-dimensional Gaussian which neglects the geometry of the trap. All three of these methods rely on a common set of conditions and assumptions I will lay out before I go into detail about them.

We make the assumption that the cloud is at thermal equilibrium and that the potential is linear where the atoms are just as we did to derive the properties of the cloud in § 5.1.1.1. The

3. The error function being defined as $\operatorname{erf}(x) \equiv \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} \mathrm{d}t$.

first assumption is justified by the fact that the trap holding time found above is approximately the thermalization time I calculate in the next section 5.2.2. The second assumption is justified by noting that the radius of average density, $R \approx 1\,\mathrm{mm}$, is much smaller than the inner radius of the MOT coils, 36 mm [1, Table 4.1], which is approximately where the potential is expected to deviate from linear (see Figure 5.1 and [170, Fig. 3]).

These assumptions let us write the potential

$$V(\boldsymbol{r}') = \kappa r', \tag{5.15}$$

where $\boldsymbol{r}' = (x, y, 2z)$ is measured from the trap center. Once the atoms have reached thermal equilibrium, their velocities will be distributed according to the Boltzmann distribution. This will read

$$p(\boldsymbol{v}) = \left(\frac{m}{2\pi k_B T}\right)^{3/2} \exp\left(-\frac{m\boldsymbol{v}^2}{2k_B T}\right). \tag{5.16}$$

where p is the probability density, $\boldsymbol{v}$ is the velocity, m the mass of an atom, k_B the Boltzmann constant, and T the temperature. When the trap is switched off, the density as given in equation (5.3) will serve as the distribution of initial positions for atoms with velocities distributed as in (5.16). The density distribution after Tof, t will then be the convolution of these two distributions in the trap.

Let the variable of convolution be $\boldsymbol{u} = (u_x, u_y, u_z) = (x, y, z/2)$ and recall that $\boldsymbol{r} = r_0 + \boldsymbol{v}t$ and $\boldsymbol{v} = \boldsymbol{r}/t$. This makes the normalized, time dependent density

$$n(\boldsymbol{r}, t) = \left(2\pi r_t^2\right)^{-3/2} \int \mathrm{d}\boldsymbol{u}\; n(\boldsymbol{u}) \exp\left(-\frac{(\boldsymbol{r} - \boldsymbol{u})^2}{2r_t^2}\right), \tag{5.17}$$

with $n(\boldsymbol{r})$ being the density distribution in the trap as in equation (5.3) and $r_t \equiv t\sqrt{k_B T/m}$. There are then different methods for extracting T from the above distribution.

For all of these methods it will by useful to define a time-dependent radius of expansion. The distribution (5.16) has a standard deviation in the velocity along any particular direction of $v_0 = \sqrt{k_B T/m}$; this will be a useful velocity scale. Let this also define a time-dependent length scale related to the expansion, $r_t = v_0 t$.

The methods that follow all start from here and use different fitting strategies to extract a temperature. Therefore they must have enough of the cloud to fit reliably and we insisted on having four standard deviations of a Gaussian fit of density profile to consider an image usable.

5.2.1.1 Variance of the time of flight

The most accurate and time consuming way we have of measuring the temperature of the cloud is the standard method of taking multiple images at different times of flight and see how the width of the cloud grows with time [1, §3.3.1]. .

When the atoms are too far apart to interact significantly (this is true both in and out of the trap, see § 5.2.2) their velocities are uncorrelated with their positions. Then the variance of their distribution in any direction can be separated into two parts, one due to the distribution in the trap and one due to the free expansion, the total variance being their sum. One part is the variance of n in equation (5.3), R^2 in the horizontal directions, x and y, and the other is

just r_t^2 defined above. This makes the variances along the two dimensions of the image plane

$$\langle (x - x_0)^2 \rangle = R^2 + r_t^2 = R^2 + \frac{k_B T}{m} t^2 \tag{5.18a}$$

$$\langle (z - z_0)^2 \rangle = \frac{R^2}{4} + r_t^2 = \frac{R^2}{4} + \frac{k_B T}{m} t^2, \tag{5.18b}$$

with $\langle (y - y_0)^2 \rangle = \langle (x - x_0)^2 \rangle$, the variance along the imaging axis.

This means that the variances in all directions scale with time like Tt^2. All that remains is to take images of thermalized clouds at different times of flight (these are assumed to be at the same temperature, making R a constant), plot the variance as a function of t^2 and do a linear fit, the temperature will be m/k_B times the slope. This means we can use the variances from a 2-D Gaussian fit of the cloud since these will scale properly with the temperature, even if the temperature cannot be reliably extracted from them because they neglect the effects of the trap. There is another limitation on this method in that we need all the atoms in the image in order to reliably extract a variance. Thus if the atoms get too hot, too few points can be taken to determine the slope of the variance with respect to t^2.

This method also, of course, involves loading, waiting in, and releasing the trap, possibly dozens of times. So while this method is the most accurate, we use it only for determining critical temperatures. It gave a typical temperature of 185 µK for cloud loaded into the magnetic trap from the MOT and allowed to thermalize for around four seconds. This forms our baseline for determining other properties of the trapped cloud, such as the collision rate.

5.2.1.2 Temperature from a single image

While our most reliable temperature estimates were obtained via the ToF method described above, we did develop a way to estimate the temperature from a singe ToF image. Here we begin from the same assumptions as above, but instead of just taking the variance of the distribution after the ToF (5.17). We do not attempt to find a two-dimensional, analytical expression to fit directly, but instead integrate (5.17) twice, once along the imaging axis and then once along another axis, to give two one-dimensional profiles that can be fit to the integrated profiles automatically calculated by our image treatment software (see Figure 2.5 for an example) to extract a temperature "along" each direction, one from the vertical profile and one from the horizontal profile. This method was used to calculate the temperature of the cloud for the data points presented in this thesis, taking the temperature associated with the more tightly confined vertical direction.

For details about this calculation, as well as the formulae for the profiles, see Appendix E.

5.2.1.3 2-D Gaussian

For day-to-day use the most important thing is comparison between clouds to make sure things are stable over the course of a working day, or during a set of runs that intend to use the same static trap data for normalization. For this, our image treatment software will fit two-dimensional Gaussian and extract a temperature based on the variance in each the two directions [1, §3.3.1.2]. This variance is fit with equations (5.18a) and (5.18b) and a temperature in each direction is extracted. These typically differ by $\gtrsim 10\%$, agreeing better the longer the gas is held in the magnetic trap prior to release. The diference being around 5% after four seconds of waiting.

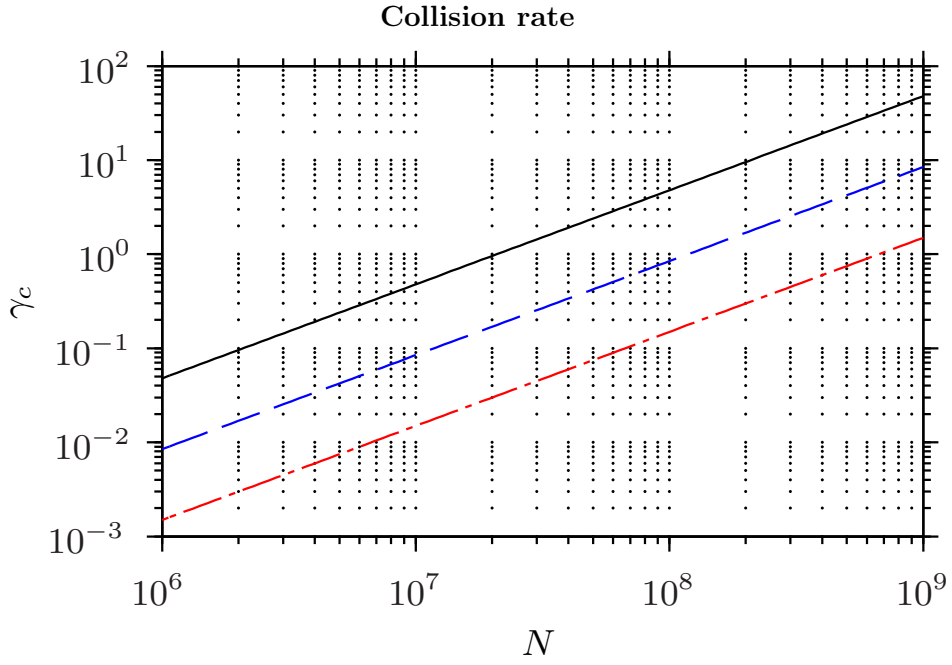

Figure 5.6 — Collision rate, γ_c, in the magnetic quadrupole trap as a function of atom number, N for three temperatures, $T = 50\,\mu\text{K}$ (black solid line), $T = 100\,\mu\text{K}$ (blue dashed line), and $T = 200\,\mu\text{K}$ (red dash-dotted line). This is for the stationary trap where $b' = 65\,\text{G}\cdot\text{cm}^{-1}$ and $\alpha = 1$.

5.2.2 Collision rate

To calculate the collision rate, we use the scattering cross-section for identical bosons [128], $\sigma = 8\pi a^2$ where a is the scattering length. The average rate of collisions, γ_c should just be this times the average relative velocity, times the average density, that is

$$\gamma_c = \sigma\langle v_r\rangle\langle n\rangle. \tag{5.19}$$

For the atoms in our trap, in state $|F = 1, m_F = -1\rangle$, $a = 2.75\,\text{nm}$ [249]. Looking back at the density (5.3) for $\alpha = 1$, as it does in the trap held in the MOT chamber, the average density is written

$$\langle n\rangle = \iiint_0^\infty n(\boldsymbol{r}')e^{-r'/R}\,\mathrm{d}^3\boldsymbol{r}' \tag{5.20a}$$

$$= C\iiint_0^\infty e^{-2r'/R}\,\mathrm{d}^3\boldsymbol{r}'\ . \tag{5.20b}$$

The integral is in equation (5.20b) is recognizable as the integral done to calculate the volume of the gas in equation (5.5) but now the integrand is squared. This factor of two in the exponent of the triple integral gives a factor of 1/8 making the average density

$$\langle n\rangle = \frac{C}{8} = \frac{N}{32\pi R^3}\,. \tag{5.21}$$

The relative velocity of two particles is derived in a similar fashion. The two velocities of the particles come from two independent Boltzmann distributions like (5.16). Their joint

velocity distribution will then be their individual distributions multiplied together, leading to a factor of two in the exponent and therefore a factor of $\sqrt{2}$ in the average, which for the distribution of a single particle is $\langle v\rangle = \sqrt{8k_BT/\pi m} = \sqrt{8/\pi}v_0$. Then the average relative velocity is written

$$\langle v_r\rangle = \sqrt{2}\langle v\rangle = \sqrt{\frac{16k_BT}{\pi m}} = \frac{4}{\sqrt{\pi}}v_0\,. \tag{5.22}$$

Putting this all together gives a scaling law for the collision rate, $\gamma_c \propto NT^{-5/2}$.

Figure 5.6 gives the collision rate as a function of atom number at several temperatures. For a typical cloud in our experiment with $N \sim 10^8$ and $T \approx 200\,\mu\text{K}$ this gives $\gamma_c \lesssim 1\,\text{s}^{-1}$. Numerical simulations [128] suggest that the thermalization time is $\gtrsim 3\gamma_c^{-1}$. This collision rate justifies our assumption the atom do not interact during a 6 ms ToF.

Unfortunately this is just about the limit of out holding time in the MOT chamber. Measuring the voltage needed to maintain current in the coil over the course of holding the atoms in the magnetic trap in the MOT chamber revealed that this voltage started to rise due to the coil heating after around 4 s. Therefore we cannot hold the atoms long enough to ensure they are thermalized, especially after a run where the atom number is likely to be lower.

Since our thermalization time can be longer than the time we can safely keep the atoms in the trap, it is worth having some way to check if the cloud is at thermal equilibrium. If the cloud is at equilibrium, then the temperature from a fit along the horizontal direction, T_x and the vertical direction, T_z, will be equal. They typically differed by $\approx 20\,\mu\text{K}$. The temperature of the run is taken along the x-axis since this tended to be in better agreement with more precise method of fitting over several times-of-flight.

5.3 Experiment and simulation

Because we are working with a quadrupole trap, we do not have access to analytic statements about the modes of the trap and have to deal with Majorana losses. Therefore we do not have precise theories about how to optimize the trajectory of the trap but can only make educated guesses and then check them by experiment. For instance, we expect large spikes in acceleration to lead to finite depth losses. We also expect trajectories that leave atoms near the center, or move many atoms quickly past the center, to lead to increased Majorana losses. We naturally thought it useful to test the typical constant velocity time profile for comparative purposes.

The experiment involved checking different time profiles at different travel times and profile parameters. We tested the standard constant velocity profile. We tested a constant acceleration profile; which would guarantee an absence of acceleration spikes. We tested a sinusoidal acceleration profile to avoid discontinuities even in the jerk. We tested profiles based on the hyperbolic tangent and on the error function, each of which offered a function smooth in all its derivatives and which could be adjusted from something like a linear profile to something like a step function via a tuning parameter, δ. We even tested an accelerating profile with a pause at $L_1/2$, detailed in D.

In this section I will focus on three of these profiles: the constant-velocity profile, the constant-acceleration profile, and the error-function profile. The constant-velocity profile is the standard strategy and therefore provides the baseline. Its naïve counter-part, the constant-acceleration profile, provides a useful comparison for understanding since it fixes the problems of the constant-velocity profile but suffers from drawbacks it does not. No profile performed

better than the error-function profile with its shape parameter, δ, going to zero, so it is given as our optimum. A few other profiles are given in appendix D for completeness.

In § 5.3.1 I will deal with the details of the experimental procedure and give the results for the three main profiles defined in § 5.1.3. § 5.3.2 will discuss the simulations we did for comparison and as a way to determine loss sources.

5.3.1 Experiment

To test a time profile the equation for the position of the trap center is first put into a Matlab script which calculates the switching times for the current and passes them along to the control software described in [1]. A run can then be taken where the magnetic trap is loaded from a MOT which had been loaded for 10s.[4] The atoms are held for a short time (200 ms) to allow them to thermalize somewhat, then moved up to coil C_7, marked in Figure 5.4, at the corner, a distance of $L_1 = 30.7\,\text{cm}$, in the transport time, t_1, and then moved back to the MOT chamber via the same profile in reverse in the same transport time. The atoms are then be held in the MOT chamber for another 200 ms before the current in the magnetic trap coils is switched off and the atoms allowed to expand freely for a 6 ms time of flight at which point they are imaged (see § 2.3.1.1). The images are then fed to the image treatment software where a standard suite of analysis is done on them giving, among other things, an atom number and the temperature, via fitting of the one-dimensional integrated profile.

These runs are typically done in sets, varying either the transport time, the shape parameter, or both. The data point at each value of the parameter is the average of three runs. In order to discount effects that are not due to the transport, such as collisions with the background gas, a set of ten normalization runs is taken for each transport time where the atoms are held in the static trap for a duration equal to that of the run, including all waiting times. We would also take a set of normalization runs after a set of data runs to make sure that the atom number had not drifted significantly over the course of data collection, which could take several hours. The average atom number of these normalization runs, $\bar{N}$ is then used as the denominator for the atom number ratio, $r_N^2 = N/\bar{N}$ for data points in that set of runs, any run in that set having atom number N. The reason the experimental ratios are given as square quantities is that we are trying to deduce the loss and heating imparted by moving the atoms once by a distance L_1. Since we must move them over that distance twice in order to measure them, the ratio we measure is the square of the ratio we want. The average one-dimensional profile fit temperature in the static trap, $\bar{T}$ is used as the denominator of the temperature ratio, $r_T^2 = T/\bar{T}$, for any individual run. Any inaccuracies due to a lack of thermal equilibrium are expected to cancel out and leave the ratio, r_T^2, a good tool for determining whether or not the transport caused any heating.

Figure 5.7 shows the effect that varying the shape parameter, δ, has on the trajectory and performance of the error function profile. It shows that the error function profile performs best when $\delta \to 0$. Whether the transport time is long or short, whether one considers heating or atom number, a lower value of δ is always better and the improvement is particularly relevant for fast transport times. Looking at 5.5 you can see that $\delta \to 0$ is where the error function profile most closely resembles the constant velocity profile, with large acceleration

4. Enough to saturate it [1, Fig. 3.14]

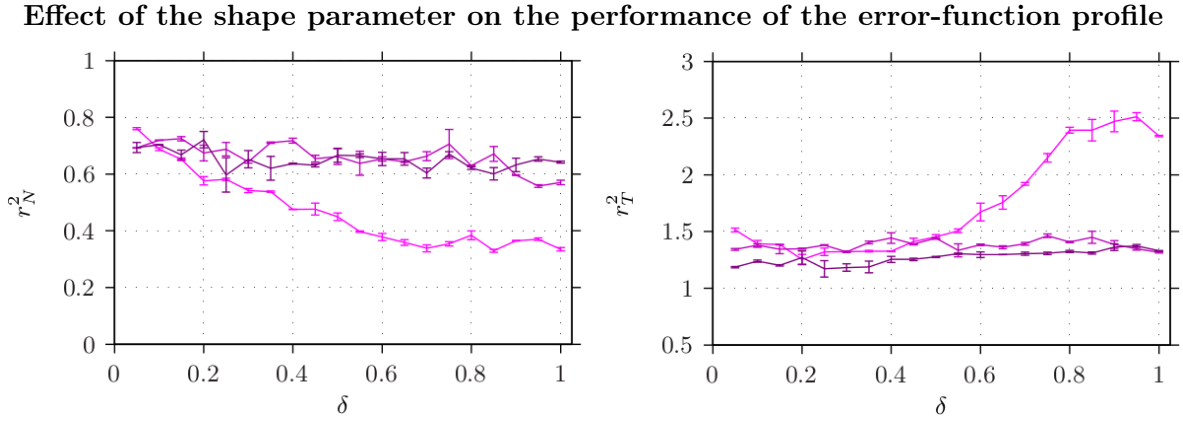

Figure 5.7 — Experimental data for the error-function profile at different value of the shape parameter, δ, as defined in equation (5.14). Setting $\delta = 0$ requires a redefinition to equation (5.23) so those values are not part of these runs. The different colors denote different travel times, Δt. From lightest to darkest they are: 500 ms, 750 ms, and 1000 ms. Left — The atom number ratio, $r_N^2 = N/\bar{N}$. Right — The temperature ratio, $r_T^2 = T/\bar{T}$.

bumps at the beginning and end. Still, it is not identical and it is written explicitly

$$y_0(t) = \begin{cases} 0 & \text{for } t \leq 0 \\ \frac{L_1}{2}\,\mathrm{erfc}\left[\ln\left(\sqrt{\frac{\Delta t - t}{t}}\right)\right] & \text{for } 0 < t < \Delta t\,. \\ L_1 & \text{for } t \geq \Delta t \end{cases} \tag{5.23}$$

This is still smooth in all its derivatives and so does not have an infinite acceleration spike, even though it is higher than for other values of the shape parameter. This is an example of experiment guiding our simulations: the optimum for the shape parameter was found experimentally and then the error function was only simulated with that shape parameter. For the remainder of this chapter, when I refer to the 'error-function profile' I am referring to the form where $\delta \to 0$ as in equation (5.23).

Data for the three profiles we are interested in is shown in Figure 5.8. The best performer is the error function profile, retaining $\approx 80\%$ of the atoms with a one-way travel time of $\Delta t = 600$ ms. No profile was able to achieve an r_N^2 value significantly higher than 0.8, even at long travel times. The standard constant velocity profile performed the worst needing one-way travel times of around 1.5 s before it achieved retention rates comparable to the other profiles. In terms of heating, all profiles which kept the maximum number of atoms also lead to heating of around 30% for a round trip, comparable to other experiments where they heated by 15% over approximately half the distance [2], and slightly under-performing [239] which heated 15% over a distance close to our round-trip length. But this seems worth the trade-off for decreasing the transport time by around a factor of two [239] or better [2], especially when a high repetition rate is called for.

The constant acceleration profile out-performed the constant-velocity profile, as expected, but it under-performed the error-function profile in terms of atom number with only a marginal improvement in heating.

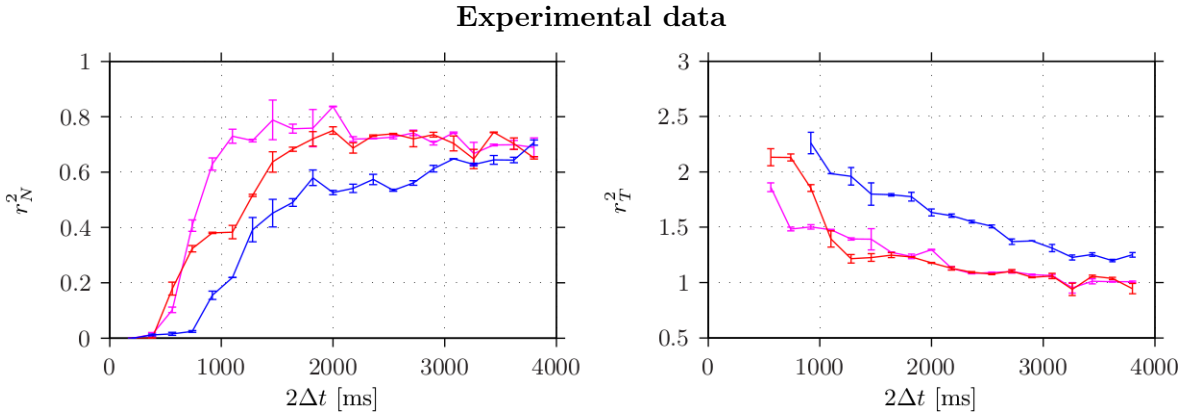

Figure 5.8 — Experimental data for the three transport profiles defined in § 5.1.3. Blue line is the constant-velocity profile, red line indicates the constant-acceleration profile and magenta the error-function profile. Left — The atom number ratio, $r_N^2 = N/\bar{N}$. Right — The temperature ratio, $r_T^2 = T/\bar{T}$.

5.3.2 Simulation

In the experiment, we can only measure how many atoms we lose, but cannot say very much about why we lose them. There are only two loss mechanisms that can be affected by moving the trap: finite depth and Majorana losses (see § 2.3.2). Since we normalize by a cloud held in a static trap, we can isolate the losses from sources associated with moving the trap but can say nothing about which mechanism predominates. This hampers our ability to make guesses about how to modify the trajectories in order move towards an optimum. Because we cannot analytically optimize our trajectories for loss rates, we made some classical simulations which will at least give us access to loss rates, both to distinguish between types of losses and make sure our overall loss rates were compatible. An overall loss rate incommensurate with our experimental data would indicate either new physics or some experimental problem. It indicated neither.

The simulations were performed in MATLAB by A. Perrin which comprised 1000 classical atoms in a infinite depth quadrupole trap with a horizontal gradient of $b' = 65\,\mathrm{G}\cdot\mathrm{cm}^{-1}$. The atoms are non-interacting to simplify calculation, and since our collision rate $\gamma_c \lesssim 1\,\mathrm{s}^{-1}$, this should not have a large effect on the outcome. The varying horizontal anisotropy was simulated from the calculations used to determine the currents in § 5.1.2. The atomic parameters are well-studied and were taken from [135]. All 1000 trajectories were calculated for a trap moving 30.7 cm in some given profile and transport time, Δt. The behavior is assumed to be symmetric for the return journey.

The fact that the trap is infinitely deep means that the potential is as written in (5.11), but equality holds for all x, y, z. This means that there cannot really be finite depth losses in these simulations. Also, since the atoms are classical and have no quantum properties, such as spin, it will not have Majorana losses. We do have their entire trajectory, including their velocity, and the value of the magnetic field at every point they visit. Therefore, instead of counting atoms lost in the simulation, we need to define some quantities that will track with losses of each type.

We expect finite depth losses to be more prevalent when the atoms are further from the trap center and that the distance most relevant to the trajectory will be along the y axis. So we look at the distance from the trap center of the i^{th} atom, $\left|\bar{y}_a^i(t)\right|$. Rather than try to find

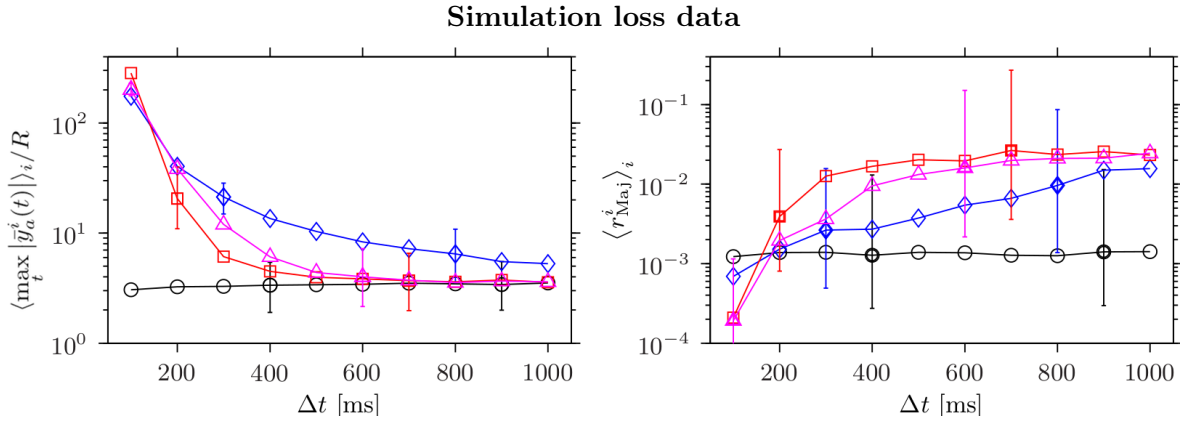

Figure 5.9 — Results of simulations with 1000 classical atoms for the constant-velocity (blue diamonds), constant-acceleration (red squares), error-function (magenta triangles) and stationary trap (black circles) for comparison. Left — Mean maximum distance from the trap minimum in units of R, proportional to finite depth losses. Right — Average maximum of the adiabatic condition for a Majorana spin-flip.

some cutoff where we call the atom "lost" we say that the loss rate should be proportional to the average of the maximum distance from the trap center, $\langle \max_t \left|\bar{y}_a^i(t)\right| \rangle_i$. This quantity, normalized by the size of the stationary trap, R, is plotted on the left in Figure 5.9. This shows that the constant velocity profile has the most finite depth losses even at $\Delta t = 1\,\mathrm{s}$. The other profiles converge with the loss rate of the static trap by $\Delta t \approx 700\,\mathrm{ms}$. Also notice that the constant acceleration profile has the least finite-depth losses, even beating the error function profile in this regard.

The other source of losses is Majorana losses which are expected near the trap center, where the magnetic field is weak. The adiabaticity condition is given by equation (2.9). When this ratio approaches unity, a spin flip becomes very likely. Since we have access to all the atomic trajectories, this ratio can be calculated for all the atoms at all times during the simulated transport. The chance of losing an atom over the course of the transport will be well-captured by the maximum value reached by this ratio. Let this be

$$r_{\mathrm{Maj}}^i \equiv \max_t \Upsilon_{\mathrm{Maj}}^i \tag{5.24}$$

where, again, i is the index of the atom and Υ_{Maj} is the adiabaticity condition defined in equation (2.9). The average of this over the atoms should give a good indication of the Majorana loss rate associated with a trajectory and is plotted to the right in Figure 5.9. Notice here that the constant acceleration profile has the most losses and constant velocity the fewest; the error function profile again falling in between them. This is expected, considering the atoms in the co-moving potential (5.8). Since the acceleration is constant they can be expected to settle near the center of the trap which to them appears stationary in a tilted potential instead of moving relative to them as with the error function profile. The same could be said of the constant velocity profile but, as we have seen, atoms in that profile tend to be further away from the trap center and so less subject to Majorana losses.

It can also be seen that for very fast travel times ($\Delta t \lesssim 100\,\mathrm{ms}$) all the profiles exhibit fewer Majorana losses than the static trap. This can be explained by them having a very high value of $\langle \max_t \left|\bar{y}_a^i(t)\right| \rangle_i / R$ and virtually no atoms being anywhere near the center. Indeed, $\langle \max_t \left|\bar{y}_a^i(t)\right| \rangle_i / R > 10^2$ means the average atom was $100R \gtrapprox 10\,\mathrm{cm}$ away from the trap

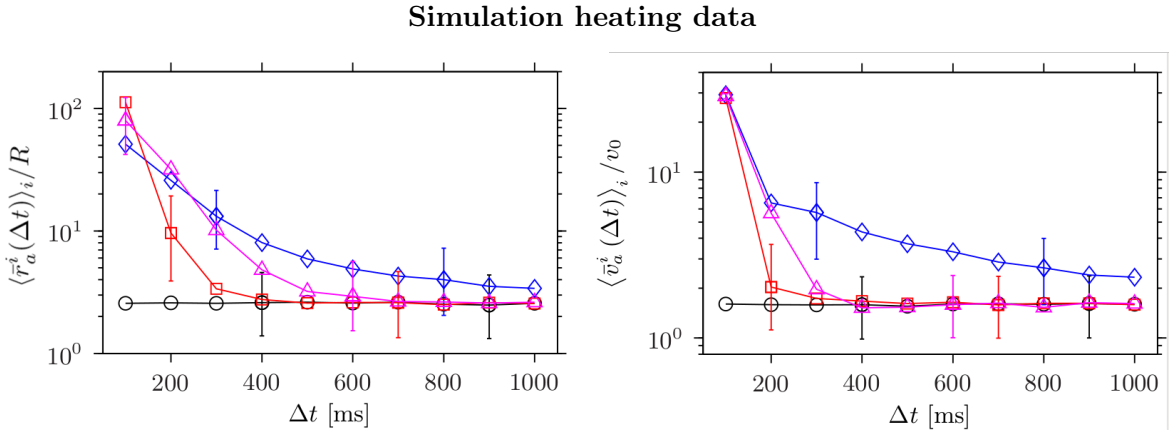

Figure 5.10 — Results of simulations with 1000 classical atoms for the constant-velocity (blue diamonds), constant-acceleration (red squares), error-function (magenta triangles) and stationary trap (black circles) for comparison. Left — Ratio of the mean distance to the trap center before and after a run. Right — Ratio of the mean of the velocity distribution after the transport to the static trap.

center, well outside the region we expect them to be trapped at all in the experiment (see Figure 5.1) and explaining the total and near-total loss of atoms at these travel times.

To deal with heating of the atoms due to transport we can look at the mean values of the position and velocity distributions. The distribution of the distances at the end of a run is $\{\bar{r}_a^i(\Delta t)\}$ where

$$\bar{r}_a^i(t) = \sqrt{(x_a^i(t))^2 + (y_a^i(t) - y_0(t))^2 + (z_a^i(t))^2}\,. \tag{5.25}$$

The final distribution for the magnitude of the velocity is then $\{\bar{v}_a^i(\Delta t)\}$ where

$$\bar{v}_a^i(t) = \sqrt{\left(\frac{\mathrm{d}x_a^i}{\mathrm{d}t}\right)^2 + \left(\frac{\mathrm{d}y_a^i}{\mathrm{d}t} - \frac{\mathrm{d}y_0}{\mathrm{d}t}\right)^2 + \left(\frac{\mathrm{d}z_a^i}{\mathrm{d}t}\right)^2}\,. \tag{5.26}$$

What is plotted in Figure 5.10 is the ratio of the average of these to the radius of average density, R defined in equation (5.4) and the standard deviation of the velocities in the static trap, v_0 defined on page 112, respectively. The higher these ratios, the more heating is expected but it will not be one even if there is no heating which is why the ratio for the static trap is plotted. Overall, they track with the experimental results.

5.4 Conclusion

Data from the experiment shows that the error-function profile (5.14) with the shape parameter $\delta \to 0$ (5.23) gives the shortest travel time while maintaining the highest atom number we observed, even at long travel times (see Figure 5.8), and gives reasonably low heating heating with $r_T^2 \approx 1.4$ at $2\Delta t = 1.2\,\mathrm{s}$, the fastest time over which we did not see significant improvement in the atom number ratio, $r_N^2 \gtrapprox 0.7$ (see Figure 5.8). This implies an optimum of $L_1/\Delta t = 51.2\,\mathrm{cm}\cdot\mathrm{s}^{-1}$ an improvement over the average speeds of $30\,\mathrm{cm}\cdot\mathrm{s}^{-1}$ [239], $12.2\,\mathrm{cm}\cdot\mathrm{s}^{-1}$ [240], or $8.25\,\mathrm{cm}\cdot\mathrm{s}^{-1}$ [2]. This is the best performing value of the shape parameter (see Figure 5.7) which makes it most closely resemble a constant velocity profile (see

Figure 5.5), the worst performer (see Figure 5.8 again) which needs travel times of several seconds to maintain the kind of atom number ratio and temperature that the other two profiles have with one-way travel time $\Delta t < 1\,\mathrm{s}$. The simulations tell us that this is due to the balance this profile strikes between the finite depth losses associated with the constant velocity profile and the Majorana losses associated with the constant acceleration profile.

While it is clear from 5.5 d.) and considerations of $\tilde{V}_d$ that profiles with large acceleration spikes need longer travel times to remain trapping ($\Delta t \approx 300\,\mathrm{ms}$) and that the constant acceleration profile avoids this problem even at fairly short times ($\Delta t \approx 150\,\mathrm{ms}$) all of these are well under the $\Delta t \approx 600\,\mathrm{ms}$ threshold at which we start to see maximum atom number. It also does not explain the performance at $\Delta t < 600\,\mathrm{ms}$ where the error function profile still wins out over the constant acceleration profile. The simulations we performed tell us that this is because while a constant acceleration profile is expected to have few finite depth losses, it accomplishes this by keeping the atoms near the center of the trap, where they undergo more Majorana losses.

The simulations tell us what proportion of the losses come from which source associated with the motion of the trap center. The constant velocity profile naturally has a lot of finite depth losses due to its high acceleration lowering the trap gradient of the trailing side significantly. When that constant velocity is high, the atoms can find themselves a significant distance from the trap center, both incurring finite depth losses and heating when the potential in the moving frame snaps upright and the atoms experience a significant increase in potential energy. The small ramp-up and down in the velocity in the error function profile seems to be enough to give the atoms enough time to catch up and prevent the large finite depth losses we see in the simulation of the constant velocity profile. The constant acceleration profile overcomes this problem by maintaining a constant, shallow tilt but this allows the atoms to settle near the center of the trap resulting in a lower average distance from the trap center (Note the consistently lower $\langle \max_t \left|\bar{r}_a^i(\Delta t)\right|\rangle_i / R$ value for the constant acceleration profile in Figure 5.10.) but that means they are more likely to enter the region of high Majorana losses and are more likely to be lost that way (see Figure 5.9). The error function profile with a low shape parameter strikes a compromise between these two loss sources leading to the best overall performance.

The heating is in qualitative agreement between the simulations and the experiment; the simulations track the comparisons between the profiles at travel times over $\Delta t = 500\,\mathrm{ms}$ but underestimate the overall heating observed in the experiment. Despite our efforts, the currents are not strictly zero when switched between sets of coils so some eddy currents are expected to lead to heating. The simulations did capture the important feature of the constant velocity profile causing a great deal of heating, even for slow travel times. This is consistent with the atoms having a higher average distance from the center of the trap, $\langle \max_t \left|\bar{y}_a^i(t)\right|\rangle_i$ under the action of that profile. This would give the atoms more potential energy which will then be converted into a higher temperature when the cloud comes to thermal equilibrium.

All of the profiles that keep the maximum ratio of atoms ($r_N^2 \approx 75\%$ for two way transport, implying $r_N \approx \sqrt{75\%} \approx 85\%$) did not have excessive heating. This may be due to our scheme of trying to maintain a constant, rather than as near to one as possible, value of α, the horizontal anisotropy. We borrow the scheme from [2] and have very similar rates of heating. For short travel times the two-way temperature ratio was $r_T^2 \lesssim 1.3$ giving a one-way ratio of $r_T \lesssim 1.14$ for a 30.7 cm trip. For a 200 µK cloud this means less than 28 µK of heating over 30 cm, right in line with the experiment in [2] in which a 125 µK cloud heated 19 µK, or 15%, over 33 cm. But where they needed 4 s to cover that distance using a constant velocity profile,

we cover a similar distance in 600 ms.

Our results show a substantial improvement over the standard constant velocity profile. We feel this is near the limit of what this guess-and-check method can accomplish without some way of identifying functions that will perform better based on their analytic properties. Although we lack analytic theory by which we can prove the error function profile is a global optimum, we hope this work will be of use to anyone transporting atoms via quadrupole trap chains.

Part II

Counters and ergodicity

Introduction

This part details the theoretical and numerical work I did during my PhD with M. Olshanii, the supervisor of my master's thesis [250]. The goal is to construct counters for the phase space of various systems that will give us insight into those systems. Let me start from the one-dimensional physics discussed in Part I and then work towards these counters, which are not a standard technique.

One-dimensional quantum systems have the possibility to be exactly solved [80], meaning analytic expressions can be found for the wave function. A good example of this would be the Lieb and Liniger solution [48] discussed in § 1.2.2.2. The condition they must meet is called *integrability*, a well defined notion in classical mechanics which means the systems has as many functionally independent integrals of the motion as it has degrees of freedom and therefore can be solved explicitly. It is not so simple to craft such a definition for quantum systems [251]. The usual meaning, in the quantum case, is independent operators which commute with each other and the Hamiltonian (e.g. [69]). This is a necessary but not sufficient condition for building an exact solution.[1] Exact solutions and questions about integrability are particularly useful in the currently active domain of out-of-equilibrium quantum systems (see [253] for a review). Exact solutions provide benchmarks that numerical methods must reproduce in the exactly solvable cases; and if a system is integrable, it is expected not to thermalize [254] or be ergodic [251].

One way of building exact solutions for systems with only contact interactions is to make a map from collisions between particles to collisions of a single particle moving in a higher-dimensional billiard [95]. Both of these will be discussed in greater detail in Chapter 6 but this is the key link between one-dimensional systems and the systems considered in this part. Via a coordinate transform I discuss in § 6.2.2 there is a straightforward map from points moving inside systems of (hyper-)planes to one-dimensional systems of particles [255]. This is the same in both the classical and quantum cases. In the quantum case, under certain conditions I mention in § 6.2.2.2, it is possible to construct exact solutions for the wave function of one-dimensional systems, like the Lieb and Liniger solution, from particular geometric arrangements of (hyper-)planes. This line of investigation is the work I contributed to during my master's thesis under M. Olshanii. Because the map from billiards to one-dimensional systems of elastic particles is exact and not merely an analogy, systemic properties such as ergodicity, chaos, integrability, and so on shown for one are inherited by the other.

1. In [80, Ch. 2], Sutherland gives [252] as an example of a quantum system proven to be integral but still lacking an exact solution.

M. Olshanii and his collaborators have also used it to solve one-dimensional subsystems and infer from their properties some properties of the system they are embedded in [256].

The idea presented here is to construct counters that will allow us to show whether or not the systems they are defined over are ergodic. This method was inspired by a 2014 paper by Wang, Casati and Prosen in [70].[2] They study a point moving in a right triangle in which the other angles are irrational and use exact numerical techniques to show that the direction the point is moving in tends to be localized near its initial direction. If this holds in the limit of infinite time, a certain class of these irrational right triangular billiards cannot be ergodic, an as yet open problem [258, Pb. 8]. Their method is to count collisions the point makes with the sides of the triangle in a particular way and then look at how this counter grows. The key link between their counter and the counter presented in Chapter 7 is that the absolute values of both grow depending on the parity of certain events in the system, not just their number.

Our work is a step towards generalizing this technique. The outline of the method is to construct counters in such a way that if they grow in a certain way then the system on which they are defined must be ergodic: either because the counter can be related to an average and the growth rate says whether or not the average will converge, or their growth is unbounded opening, the way for a measure-theoretic argument. Then one finds the growth rate of the counter using cycles of growth and stasis that will appear in all parity counters when there are periods in the dynamics of the underlying system .We start with a counter defined over a textbook example of a simple ergodic system, an irrational rotation, and use that as a way to develop tools for looking at the growth of counters that are periodic on short time scales. That is, the sequence given by the counter can be replicated by periodic subsequences for certain lengths of time. This itself takes inspiration from a 2009 paper by Huveneers [259] where he constructs unbounded subsequences in irrational rotations by looking at subsequences of rational approximants. We give a way of calculating this range of validity for the periodic approximations and I will discuss what this means for where the particular methods we develop can be expected to fail.

With the goal of introducing new tools which will help us gain insight into problems of physical significance established, I will lay out the background theory I will be relying on in Chapter 6 before discussing the inspirational problem and the work done on the method in Chapter 7.

2. I will at times make reference to their 2013 preprint of the same article [257] since it contains some numerical results that the 2014 published version [70] does not since the argument differs between the two.

Chapter 6

Background

In this chapter I will be laying out the concepts necessary to understand the next chapter which details the results for Part II. The goal of this research is to construct objects we call counters for dynamical systems, examine the behavior of those counters, and infer from that behavior ergodic theoretic properties of the system. While we take the basic idea from Wang, Casati and Prosen [70] (I use the acronym WCP to indicate objects and systems coming from their paper), who defined their counter over a billiard with clear physical interpretation. They were not able to say with certainty what the behavior of their counter implied. For the counter we study in detail, we define a counter over a simpler system, proven to be ergodic but with less obvious physical significance, the irrational rotation. We use this as a test case to develop tools for analyzing the growth of counters that I will go into in the next chapter.

If I am going to do that, I had better explain what all of those things are. I will also make brief digressions where appropriate to both include small points that will be relevant in the next chapter and to highlight connections to the physics discussed in Part I. I will start by laying out the ideas I will need from ergodic theory, including some basic ideas from measure theory that will be needed. I will mention why a system having memory (i.e. being sensitive to its history) is incomparable with it being ergodic since Wang et al. want to argue from the memory they seem to have found in their counter to the non-ergodicity of the billiard. I will then define a billiard in detail and show how they can encode one-dimensional systems. I will also briefly explain how they can be used to build exact solutions to quantum systems, like those studied in Part I. Then I will discuss rotations, both rational and irrational as well a few results from the theory of continued fractions to deal with how to generate series of rational approximants.

6.1 Ergodicity

6.1.1 Definitions

The fundamental question of ergodicity is whether or not infinite time averages converge to ensemble averages. This is very useful because one of these averages can be much easier to find than the other. If one can then show that some system or transformation is ergodic and find one limit, one has the other immediately. As mentioned in the introduction, physical interest goes back to Boltzmann [58] who wanted to show this relationship between averages

for an ideal gas if he could prove that in infinite time. It has been shown that his specific notion of ergodicity is impossible under Newton's laws [60,61] but that his notion of a trajectory densely filling the phase space after infinite time, called *quasi-ergodicity* is sufficient to show the convergence of the averages [63]. I will not make such fine distinctions as the one between quasi-ergodicity and another kind and will simply use the term *ergodic* for systems and observables whose infinite time and ensemble averages converge.

I will generally be dealing with ergodicity in measure theoretic terms (ergodic theory being known also as *measurable dynamics* [3, p. 1]) and [3] provides a good reference for the ideas I will collect here. The appendices of [260] also provide a rigorous but much condensed introduction. More general references and applications are almost too numerous to list but [261] along with those already listed provided my basis in the literature.

Let me define some terms before I briefly return to the physical significance at the end of this section since it will make the explanations less circuitous. I will use the term *system* to mean a phase space, comprising a set X containing phase points or *states* of the system, $\{x\}$, and having a σ-algebra,[1] Σ, and being equipped with a map, $T : X \to X$, that updates the states. An *observable* of this system is a function of its states. Both systems and observables can be said to be ergodic and if a system is ergodic then all of its observables are, but non-ergodic systems can have ergodic observables [262].

A *measure* is a map $\mu : \Sigma \to \mathbb{R}^+ \cup +\{\infty\}$ such that for disjoint sets $E_i \in \Sigma$, $i = 1, 2, \ldots$

$$\mu(E_i) \geq 0 \quad \forall\, i \tag{6.1a}$$

$$\mu\left(\emptyset\right) = 0 \tag{6.1b}$$

$$\mu\left(\bigcup_{i=1}^{\infty} E_i\right) = \sum_{i=1}^{\infty} \mu\left(E_i\right). \tag{6.1c}$$

These are all the desired properties for something which generalizes the idea of size. We want it to give a number in $\mathbb{R}$, which is positive when it measures something (6.1a) and zero when it measures something with no size (6.1b). Lastly, the size of a collection of things should be the sum of their sizes when measured individually (6.1c). If a measure is defined only on a certain set, that set is said to have *full measure* and if that set is finite, it can be said to have *measure one* and the measure to be a *probability measure*. A set of full measure can be said to cover *almost everywhere* and contain *almost all* points. The word "almost" in both of these expressions conveying the sense of "except for a set of measure zero." If a transformation keeps all of sets in Σ the same size as measured by μ, then T is said to be *measure preserving* or equivalently that μ is an *invariant measure* for T or a *T-invariant measure*. The explicit condition that must hold for this is that all measurable sets $E \in \Sigma$ must be mapped *from* sets of equal size by T, i.e. $\mu(T^{-1}E) = \mu(E)$.

This is first time I want to make a link back to ergodic theory explicit. When doing the ensemble average, one can imagine picking each point in the phase space once and only once, calculating the observable there, and adding it into the average. Therefore the odds that states in some region of phase space will affect the ensemble average only depend on the size that region takes up. This makes the expression "probability measure" clear, the probability that a point is anywhere in the phase space is one, and this leads naturally to concept of the

1. A σ-algebra on X is a collection of subsets, including $\emptyset$, of X, along with their compliments (if $A \in \Sigma$ then $A^c = X \setminus A \in \Sigma$) and countable unions (if $A_i \in \Sigma$ for $i \geq 1$ then $\cup_{i=1}^{\infty} A_i \in \Sigma$). This is taken from [260] from which I take my definition of measure. For this thesis, you can think of Σ as the countable subsets of X, including itself, as distinct from the points in X.

ensemble average as the expectation value of an observable $f(x)$, weighted by the measure of the volume element phase space, which reads

$$\langle f \rangle \equiv \int f(x)\, \mathrm{d}\mu\,, \tag{6.2}$$

where $\mathrm{d}\mu = \mu(\mathrm{d}x)$. This is one part of what it means to be ergodic, to deal with the other part we need to deal with how point in the phase space move with time.

The map that is the solution to a set of time-independent[2] differential equations and is a *flow*. A flow can be continuous to discrete. I use the discrete notation here since the systems treated in the next chapter progress by well-defined time steps, but in either case a flow is a transformation that, when acting on X, has the properties [263]

$$T^0 x = x \tag{6.3a}$$

$$T^t(T^s x) = T^{t+s} x\,. \tag{6.3b}$$

Here t and s can be thought of as lengths of time for a continuous system or a number of units of time in the discrete case. Then (6.3a) means that if no time passes, the system remains in the same state and (6.3b) means that the result of a transformation by a flow depends only on the amount of time which passes since the system was in the state, x. The action itself is always the same sine the differential equations it solves are time-independent. If we take $x' = T^s x$ then there is now way to distinguish a system that began at x' a time t ago, or began at x a time $t + s$ ago; regardless of the history the state and the next state, after another application of T, are the same. In this sense, a flow has no "memory," an idea I return to in § 6.1.2.

Where this gains so much physical significance is that these time-independent differential equations are just what govern isolated systems in classical mechanics. Not only are these phase spaces measurable, but as a result of Liouville's theorem [264, §46] we can say that such flows are measure preserving and regions of the phase space will be mapped onto regions of identical size no matter how long the system evolves for, as long as the flow is constant. If a particular region of the phase, a set $A \subseteq X$ for instance, is mapped back onto itself, that set is said to be an *invariant set*. Note that points need not return to themselves, it is only required that points in A never leave A under the action of T.

According to Birkhoff's ergodic theorem [64], a system is ergodic[3] if and only if all of the invariant sets, $A \in \Sigma$, of a system have the property that $\mu(A) = 0$ or $\mu(A^c) = 0$ [3, Lm. 3.7.1]. This property is called being *metrically transitive.*. That is to say that all the invariant sets in the phase space have either full measure or measure zero. For a rigorous proof see the references, but to get an intuition that this implies ergodicity, imagine an invariant set $A \subset X$ has positive but not full measure. Then all trajectories that start in A will never visit any states in A^c. Then the partial time average

$$\overline{f}_n(x_0) \equiv \frac{1}{n} \sum_{k=0}^{n-1} f(T^k x_0)\,, \tag{6.4}$$

will only converge to the ensemble average (6.2) by some extreme coincidence for $x_0 \in A$, since the ensemble average will also include all of the points in A^c which has measure $\mu(A^c) =$

2. That is that the only time dependence enters through the variables, the equations themselves having no explicit time dependence.
3. To be exact it is quasi-ergodic, his 1931 paper being a follow up on von Neumman's paper [63] which was published a few months later in 1932.

$\mu(X \setminus A) > 0$. If, on the other hand, A is the only invariant set and has full measure (i.e. $\mu(A) = \mu(X)$), then the flow will take all of the trajectories throughout the phase space. The flow cannot take a trajectory back to any individual point since by property (6.3b) it must then visit all the same points again, forms a closed curve, and we have found an invariant set with less than full measure, since a trajectory is of one-dimension and therefore cannot fill a higher-dimensional phase space, as mentioned above. It also cannot get stuck in any particular region of the phase space since, again, A is the only invariant set and it has full measure. Then, given infinite time the trajectory must densely cover the phase space and this is enough to ensure that the infinite time average

$$\overline{f}(x) \equiv \lim_{n\to\infty} \overline{f}_n(x)\,, \tag{6.5}$$

converges to the ensemble average (6.2) and ergodicity is satisfied [62, 64]. It just remains to notice that this argument does not break down if there are fixed points in the phase space where trajectories will never leave. Since they are points, they have measure zero and therefore will have no weight in the ensemble average and the probability that a trajectory will reach them in the time average is zero.

At infinite time, when the phase space has been densely covered, there can be no notion of one region of phase space being "more densely" covered than another. When we say that a trajectory beginning at x_0 covers the phase space *densely* we mean that for any point in the phase space, y, and a distance $\epsilon > 0$, there is a t such that the trajectory $T^t x_0$ lies within ϵ of y, no matter how small ϵ is chosen. Since $t \to \infty$ in (6.4) there is no sense in talking about regions where t need not be as large for a given ϵ or vice versa.

6.1.2 Systems with memory

I mentioned before that equation (6.3b) means that flows poses no "memory" and I would like to come back to that, now that we have some helpful concepts defined. When I say that a system has *memory* I mean that a point on a trajectory itself contains no information about that trajectory. This includes any sensitivity to the history of the trajectory, not just the case where the entire history can be constructed from a phase point. The flows themselves clearly do not posses this property, owing to the fact that the differential equations they solve do not depend explicitly on time, but that does not mean all systems evolving under flows also lack memory.

An example would be a simple harmonic oscillator of frequency ω and mass of one. Given a point in its phase space, (q, p), where q is position and p is momentum, the trajectory could only ever come from an ellipse consistent with its energy, $\mathcal{E}$, and therefore can be said to have memory of $\mathcal{E}$. This could be used to define an invariant set of points, $\{p, q\}$, satisfying the relation

$$p^2 + \omega^2 q^2 = 2\mathcal{E}. \tag{6.6}$$

But this has measure zero in the phase space of the entire q, p plane and so it is not obvious that this system is non-ergodic given what we have said before. But we can also define an invariant set of all the points lying in a strip between the ellipses

$$p^2 + \omega^2 q^2 = 2\mathcal{E} \pm \epsilon \tag{6.7}$$

for some $\epsilon > 0$ and p, q will have memory as coming from this strip. This strip will be in the σ-algebra of the space, and have positive, but not full, measure.

This gives a model for how non-ergodic systems can posses a memory. If there is at least one invariant set of measure strictly between zero and one, then states in those invariant sets can be associated with an invariant of the motion and so always keep a trace of where they came from. States not in invariant sets will be characterized by a lack of any of these constants of the motion and therefore show they came from outside the invariant sets. The converse case, that a system with memory is not ergodic, is only true if the memory persists for infinite time and is inherent in the phase variables. A flow may map a subset of positive measure to itself for a long but finite time, for instance, but the system still contain no strictly invariant sets. If one constructs some object that keeps track of the time passed as well as the position in phase space, that will have memory of the entire trajectory, but only because it was so constructed. But persistent memory in a system is strong, but not conclusive, evidence that it is not ergodic.

KAM theory The way invariant sets in phase spaces are usually found by physicists is with Kolmogorov-Arnol'd-Moser (KAM) theory. The most basic statement of the result of the theory is that phase trajectories of many systems with time-independent Hamiltonians will be confined to invariant tori of positive measure and it describes how to find these tori. More precisely the system must be *integrable*, that is it must possess as many independent constants of the motion as degrees of freedom. For more the reader is directed to [265] excellent for its review of the literature if nothing else. It also shows that unless the trajectory is periodic (i.e. belongs to an invariant subset of the torus) it densely covers the torus. It also shows that these tori persist under small perturbations in the Hamiltonian away from integrability but that is getting away from my reason for bringing it up. The relevant consequence is that unless there is only a single such torus, exactly coinciding with the entire phase space, the system is not ergodic. Given a system that this theory applies to, it would be surprising for the system to be non ergodic for some other reason. This is what Wang et al. have found numerical evidence for in their paper on triangular billiards [70].

6.2 Billiards

Now that I have outlined the kinds of questions we will be asking, I will go over the systems which inspired the investigation presented here and provide the link to physical, one-dimensional systems.

6.2.1 Introduction

Mathematical billiards are a way of considering fundamental questions in dynamical systems (a much broader class of problems [4]), consist of a point, moving in straight lines inside some domain, the boundaries of which it bounces elastically off of [66]. Since billiards can encode one-dimensional systems which can include objects of infinite mass, I will refer to the infinite-mass objects in one-dimensional systems as *walls* and I will say that the boundaries of billiards comprise *mirrors*. In the most basic case, the domain is a flat, two-dimensional polygon [266] but this can be extended to domains with elliptical boundaries (where the billiards are integrable [267]), piece-wise curved boundaries which can thermalize trajectories [268] as well as in other geometries and in the presence of fields [269]. It is possible to define *outer* or *dual* billiards [270] where instead of a point reflecting inside a closed boundary, a point outside a boundary evolves step-wise by being reflected along lines tangent to that

boundary which can be used to study celestial mechanics [271]. See [272] for a survey. The textbook [66] provides a good introduction and overview on the subject and has an extensive bibliography including several reviews.

The questions we will examine with respect to billiards are those of ergodicity and periodicity. The trajectory of a billiard is periodic if it returns exactly to some point in the phase space; that is it returns to the same position and the same velocity direction. Since the future trajectory at any moment depends only on the point in phase space occupied at that moment (that is to say that billiards have no memory and motion in them constitutes a flow as defined in equations in(6.3a) and (6.3b)), the first time the trajectory revisits a point in phase space it will fall into an *orbit* and revisit all prior points in phase space and no others. These orbits constitute sets which are invariant under the flow of the billiard motion. If these periodic orbits have full or zero measure then the billiard is ergodic. Trajectories which meet the corners tend to be ignored unless there is a clear way to continue them, typically with some physical justification [273]. But these trajectories form a set measure zero since the corners themselves have measure zero.

A billiard or its boundary is said to be *rational* or if all of the angles in its boundary are rational multiples of π. These are referred to as *rational angles*. Otherwise, the billiard is *irrational*. Since we will be dealing with right-triangular billiards and examining the behavior when the other angles are nearer or further from rationality, it is significant that billiards in near-rational polygons have been shown to be ergodic [274]. For right triangles far from rationality the question is much more open. It has been shown that almost all trajectories perpendicular to the legs of an arbitrary right triangle are periodic [275]; that other types of periodic trajectories exist in right triangles and that every point on the triangle is on at least one such trajectory [276]; and that any periodic billiard path in a right triangle is unstable under perturbations of the angles in the triangle [277].

The question Wang et al. were approaching in [70], is one of the ergodicity of trajectories in irrational right triangles when the initial position and velocity direction can be anything, an open question [272, Plbm. 8]. The basic strategy is construct something that counts the distinct velocity directions the trajectory visits. For a right triangle with irrational angle $\alpha/2$, each reflection off of the hypotenuse will change the velocity direction by 2α (see Figure 7.1) so the set of potential velocity directions is $\{2\alpha\mathbb{N}\}$, as set of cardinality $\aleph_0$. A set of trajectories of positive measure for which the number of velocity angles visited was finite would therefore indicate that these triangles are not ergodic since not all of the accessible angles are visited. This in turn would tell us about a system of two masses colliding in a one-dimensional box.

6.2.2 Encoding one-dimensional systems

The physical interest in billiards for us is that maps can be made between billiards and one-dimensional systems [255]. There are different ways of making the encoding, but the general idea is that given a system of N particles with N positions in one-dimension, one can transform coordinates and follow a single *virtual particle* in N dimensions.[4] The mirrors of the billiard then correspond to collisions between the particles. This correspondence is illustrated in Figure 6.1.

For classical, impenetrable, particles, the relevant billiards are just those discussed above: the virtual particle moves in straight lines and reflects elastically off the mirrors. That is the kind of billiard I will be working with in this part and I will discuss it first and go over the

4. Or $N-1$ dimensions, if one is only interested in the relative motion of the particles.

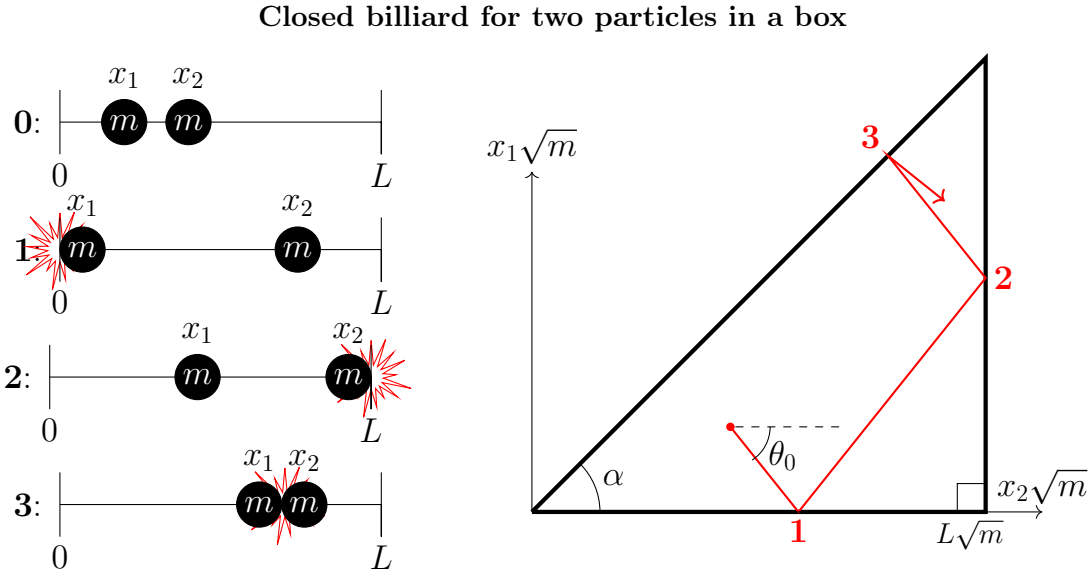

Figure 6.1 — Illustration of the correspondence between two particles in a box (left) and a billiard (right). Because the masses are equal, the coordinate transform outlined at the beginning of § 6.2.2.1 implies that the billiard is a $45° - 45° - 90°$ triangle. Collisions in both systems are labeled 1 to 3. The initial particle velocities (not shown) are encoded in $\theta_0 = \arctan(\tilde{v}_1/\tilde{v}_2)$ (see text).

coordinate transform. Similar maps can be made in the quantum case, however, and exact solutions such as the Lieb–Liniger model [95] and systems of mass mixtures (see [69] and references therein) can be constructed. I touch on them to draw a more explicit connection between billiards and the systems discussed in Part I.

6.2.2.1 Coordinate transforms and virtual particles

Billiards can be directly related to the dynamics of one-dimensional systems of hard-core particles via a coordinate transform [66, Ex. 1.2]. Let x_i, with $i = 1, 2$ be the positions of two impenetrable particles indexed by i of mass m_i moving in one dimension with velocities v_i. By scaling the coordinates like $\tilde{x}_i = x_i\sqrt{m_i}$ the velocities also rescale, becoming $\tilde{v}_i = v_i\sqrt{m_i}$. An example where $m_1 = m_2 = m$ and the particles are confined to a box of length L is illustrated in Figure 6.1. The place where the particles collide, $x_1 = x_2$ becomes a line in the transformed space, making an angle $\alpha = \operatorname{atan}\sqrt{m_1/m_2}$ to the $\tilde{x}_2$-axis. This line is colinear with the vector $(\sqrt{m_1}, \sqrt{m_2})$.

Conservation of energy and momentum become

$$m_1 v_1^2 + m_2 v_2^2 = m_1 v_1'^2 + m_2 v_2'^2 \to \tilde{v}_1^2 + \tilde{v}_2^2 = \tilde{v}_1'^2 + \tilde{v}_2'^2 \tag{6.8a}$$

$$m_1 v_1 + m_2 v_2 = m_1 v_1' + m_2 v_2' \to \tilde{v}_1\sqrt{m_1} + \tilde{v}_2\sqrt{m_2} = \tilde{v}_1'\sqrt{m_1} + \tilde{v}_2'\sqrt{m_2} \tag{6.8b}$$

where primed quantities are those after the collision. If we let the velocity vector of the virtual particle be $\boldsymbol{v} = (\tilde{v}_1, \tilde{v}_2)$ then (6.8a) says that the magnitude of $\boldsymbol{v} \cdot \boldsymbol{v}$ does not change when the particles collide and (6.8b) says that $\boldsymbol{v} \cdot (\sqrt{m_1}, \sqrt{m_2})$ does not change when the particles collide. This is just to say that the virtual particle reflects off of the $x_1 = x_2$ line just as in a billiard (**3** in Figure 6.1). When $x_1 = 0$ the first particle meets the wall at $x = 0$ and $v_1 \to -v_1$ so the component of $\boldsymbol{v}$ in the $\tilde{x}_1$ direction changes sign (**1** in Figure 6.1) and for a

Method of images on an open billiard for two masses on a half-line

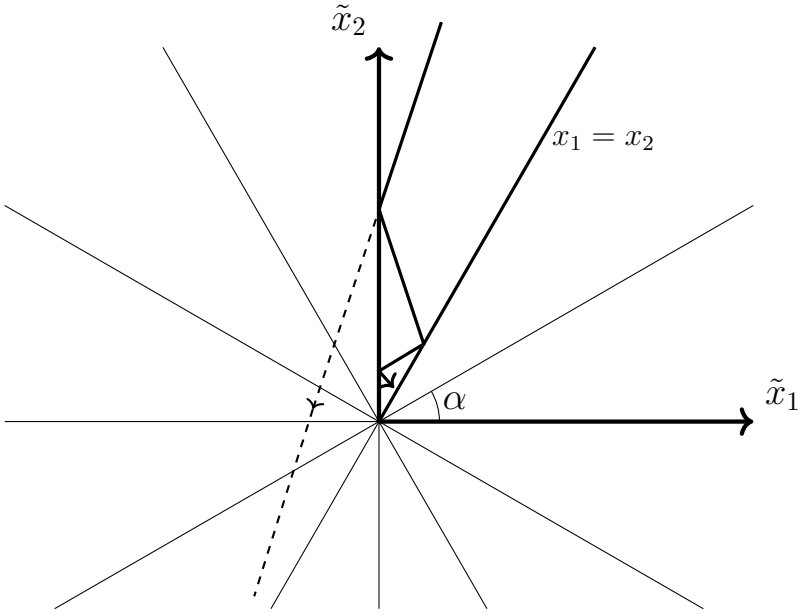

Figure 6.2 — Billiard encoding two particles of mass $m_{1,2}$ at $x_{1,2}$ moving on the half-line. The solid black arrow depicts the billiard trajectory. The dashed arrow represents the *unfolded* trajectory (see text) and thin solid lines the unfolded mirrors or *mirror images*. Since α equals π over an integer (6 here) the mirrors reflect each other onto each other and form a closed group [278] and the entire set generated by reflections can be, and is, depicted.

box of length L the analogous thing occurs at $x_2 = L, \tilde{x}_2 = L\sqrt{m}$ (**2** in Figure 6.1). Thus we have exactly the properties of a point moving in a closed billiard.

A few things are worth nothing before we move on to the quantum case. If the particles are hard-cored, the virtual particle will never meet the $\tilde{x}_1$ axis away from the origin as this would mean $x_2 = 0$ while $x_1 > 0$; the particles have changed order. The other thing to notice is that instead of reflecting the virtual particle, one can reflect the billiard itself about the mirror the virtual particle is colliding. This *unfolds* the path of the virtual particle into a straight line across many copies of the billiard. This is sometimes called the *method of images*. It goes back to electrostatics where charges induce *image charges* in conductors [279, pp. 185 and following] but has since been generalized to Markov processes and solving certain differential equations with Dirichlet boundary conditions [280, 281], which includes flows on billiards. This will be useful later for seeing the connection between billiards and rotations.

Take the example illustrated in Figure 6.2, the open billiard encoding a system of two, hardcore particles on a half-line. The virtual particle enters at an angle equal to $\text{atan}(3) \approx 2\pi/5$. Following its path via the method of images, it will keep making collisions until the next mirror image it would meet with is at a steeper angle; the $\tilde{x}_2$ axis in this case. Thus it will make six collisions before it emerges from the open billiard. The order in which those happen is given by the trajectory of the virtual particle. First, $x_1 = 0$ and that particle hits the wall, then it hits the other particle, then the wall, etc. The images encode the collisions associated with the mirrors they are images *of*.

For studying systems where only the relative motion of the particles is of interest, one can transform to coordinates [282] similar to the Jacobi coordinates; the first particle is said to be

at x_1 and then the position of particle $i > 1$ is measured relative to the center of mass of all the particles $j < i$. Then, one of the dimensions of the billiard is associated with the center of mass position and this dimension is orthogonal to all the others and can be discounted.

6.2.2.2 Quantum billiards and the Bethe ansatz

While many interesting properties of quantum systems can be investigated using billiards [283, 284], with some extension of the analogy, exact solutions to one-dimensional systems of δ-interacting particles, can be constructed from billiard systems. The methods go back to McGuire solving several cases in 1963 [87] and whose optical analogy I will summarize here. Gaudin introduced the idea of working from billiards to physical systems in the 70s [95], using it to reconstruct Lieb and Liniger's solution. This work has been continued by M. Olshanii, Harshman, Jackson, and their collaborators [69, 99, 285], including myself [250].

The solutions rely on the coordinate Bethe ansatz [68, 97],[5] just as Lieb and Liniger did [48]. The ansatz consists of assuming that the solution is a sum of plane waves, and is symmetric (in the case of bosons) or anti-symmetric (in the case of fermions) under exchange of particles. This (anti-)symmetry lets one turn one true equation into a system of equations large enough to constrain the variables and construct an exact solution. The way to do this with billiards is to replace the initial virtual particle with an initial plane wave; this constitutes a *quantum billiard* [286]. The reflections then correspond to the permutations of the particles and the solution is the sum of the initial wave and of all its reflections [87]. In the case where reflections about the mirrors form finite reflection groups [67] these can be solved exactly using the method of images [98]. One can then start from such a set of mirrors, and construct a wave function. Then the difficulty is assigning physical significance to the mirrors. If one mirror is reflected onto another, they must encode the same interaction; if the particles are not hard-cored, the wave at the mirror will be partially transmitted and the method of images may fail; the angles between the mirrors cannot suggest negative masses; etc. By doing things like embedding the billiard in a curved manifold, one can add external potentials to the system [69], and by finding new mass ranges that can be related to finite reflection groups [99, 250], the set of exact solutions of one-dimensional systems can be expanded. Thus billiards remain an active area of research relevant to the systems discussed in Part I.

Let me describe a simple example, the three impenetrable particles on an open line treated by McGuire [87]. The system has three masses so it needs two degrees of freedom, one for each of the relative velocities. After the coordinate transform [282] this gives a two-dimensional, open billiard, similar to the one illustrated in Figure 6.2. But now, instead of the incoming arrow depicting a billiard trajectory, imagine it is the wave vector of an incoming plane wave. Since this is the sum of two plane waves along each axis, it alone encodes the two relative states of the three particles. What the Bethe ansatz relies on is summing over all the plane waves that solve the differential equations of the system.

We have two of them already in the initial wave vector, where do we get the rest? The Bethe ansatz says to sum over the set of permutations of the particles. This set can be generated by swapping neighbors, which in momentum space is equivalent to neighbors undergoing elastic collisions [48].

In a closed system, like the Lieb and Liniger system discussed in § 1.2.2.2, this means summing over all the permutations and the billiard boundary must be a closed curve. When the system is open and the billiard boundary is just a wedge, it means summing over the

5. As distinct from the algebraic Bethe ansatz

collisions the system will actually undergo. Those are just the collisions encoded by the billiard mirrors the incoming wave reflects off of. Therefore, the set of plane waves to sum over is that generated by a plane wave under the action of a group of billiard mirrors. The wave front emerging after all the reflections have taken place will be the sum of all the permutations the system will undergo and so it, plus the incident wave, provides the solution. If the reflection group of the mirrors is finite[6] then this has no diffraction effects, there is a single emerging wavefront which evenly fills the whole wedge and the sum is readily computed. If not, there are extra terms in the sum which resist exact solution (see [87, §III.A] and references therein).

For examples of one-dimensional systems solved with billiards and reflection groups, see [95], [287], or [285] and references therein.

6.3 Irrational rotations

The counter studied in the next chapter is inspired by one defined on a billiard but it itself is defined over a *rotation*. The phase space of a rotation is a circle and the flow moves phase points counter-clockwise by some angle, α. The rotation is said to be *rational* or *irrational* depending on whether α is rational or irrational. Irrational rotations are such basic ergodic systems that they are used in textbooks to introduce the idea [3, §3.2], [261, Prop. 2.16]. They are well-studied objects in mathematics [259,288,289], and have straightforward physical relevance via their close connection to the almost-Mathieu operator [290]. This is a discrete, one-dimensional operator defined on square integrable functions which is closely related to the Hamiltonian of electrons in low-dimensional lattices subject to a magnetic field (see [291] and references therein). Our interest in them though lies in using them to gain insight into the properties of billiards. Because a billiard trajectory can be unfolded via the method of images, and reflecting about two intersecting lines (or planes) is equivalent to a rotation by twice the angle between them [292, p. 33] rotations can play an important role in understanding billiards [293, §7.5]. Since the WCP counter has interesting implications about the behavior of billiards, our strategy is construct a related counter over this simpler system which we know to be ergodic, study its behavior, and then, in the future, try to link it back to the billiard problem.

In the next chapter, the methods for studying these counters rely on examining how their rational convergents behave; so a few details about rational approximations and some specifics of rational rotations will be necessary. I will then cover the notions we will need from irrational rotations themselves.

6.3.1 Rational approximants

Any irrational number can be represented by an infinite continued fraction (see [294, Thrm. 170]. I take my basic facts about continued fractions from this text unless otherwise noted.). This is constructed from an infinite series of integers, a_i, $i = 1, 2 \ldots$ with $a_1 \in \mathbb{Z}$ forming the whole part and $a_i \in \mathbb{N}^+$ for $i > 1$. Note that this is slightly idiosyncratic (but not original [295]) notation, i typically beginning at zero and a_0 being the whole part of the number. The

6. A complete list of finite reflection groups has been given by H. Coxeter [67]. It requires that the angles between the mirrors be certain integer ratios of π.

continued fraction representation of the irrational is constructed from this sequence thus

$$\sigma = [a_1;\, a_2,\, a_3,\, a_4,\, \ldots] = a_1 + \cfrac{1}{a_2 + \cfrac{1}{a_3 + \cfrac{1}{a_4 + \ldots}}}. \tag{6.9}$$

I will call σ a *generic irrational* if the corresponding sequence, $a_1, a_2, a_3, \ldots$, is not bounded from above. These constitute almost all irrational numbers [294, Trm. 196], hence the name. That is, if σ is a generic irrational then

$$\sigma = [a_1;\, a_2,\, a_3,\, a_4,\, \ldots]: \quad \forall A \in \mathbb{N} \quad \exists\, i : a_i \geq A. \tag{6.10}$$

I will call a σ *constant type* (after [259]) or *bounded irrational* if the corresponding sequence, $a_1, a_2, a_3, \ldots$, *is* bounded from above. That is, if σ is of constant type

$$\sigma = [a_1;\, a_2,\, a_3,\, a_4,\, \ldots]: \quad \exists\, A \in \mathbb{N} : \forall i\; a_i \leq A. \tag{6.11}$$

The rationale for calling irrationals with bounded continued fraction series "bounded" is clear. The fact that the set of such irrationals has measure zero [296, §7] justifies the calling ones that are not bounded "generic." If an irrational number is of generic type, it can be approximated arbitrarily well since large numbers in the continued fraction series imply a better approximation by the truncated fraction and somewhere in the continued fraction series can be found numbers of arbitrary magnitude.

The distinction between these two types of irrationals will be significant for estimating the growth of the counter introduced in Section 7.2. The methods we use fail for irrationals of constant type because we cannot be sure if the behavior we see in the approximate series holds long enough to give an estimate of true series.

Rational approximations of σ are generally written

$$\sigma_M = [a_1; a_2,\, a_3,\, \ldots,\, a_M] = a_1 + \cfrac{1}{a_2 + \cfrac{1}{a_3 + \cfrac{1}{\ddots + \frac{1}{a_M}}}} = \frac{p_M}{q_M}\,, \tag{6.12}$$

In the series of rational approximants to a given σ, it follows from their construction via the Euclidian algorithm [294, § 10.6] that the denominator, q_M, grows with M like

$$\frac{q_M}{q_{M-1}} = a_M + \frac{q_{M-2}}{q_{M-1}} \quad \text{for} \quad M \geq 3\,. \tag{6.13}$$

The first two denominators will always be $q_1 = 1, q_2 = a_2$. From (6.13) it follows that for generic σ where a_M can grow arbitrarily large,

$$\frac{q_M}{q_{M-1}} \overset{M \gg 1}{\approx} a_M\,, \tag{6.14}$$

holds very well. Indeed, in the generic case $a_M \overset{M \to \infty}{\to} \infty$, but for any continued fraction, one is guaranteed that $\frac{q_{M-2}}{q_{M-1}} \to 0$ as $M \to \infty$ so it can be neglected from (6.13) and the

approximation (6.14) can be expected to be good for large M and arbitrarily well in the case of generic σ.[7]

By putting together some basic theorems from the theory of continued fractions, one can show that

$$\frac{1}{q_M(q_{M+1}+q_M)} \leq |\sigma_M - \sigma| \leq \frac{1}{q_M^2 a_{M+1}}\,. \tag{6.15}$$

The important thing to notice is that, in the case of generic σ, the rightmost term can be made arbitrarily small and the approximation arbitrarily better than $1/q_M^2$, a distance from the irrational any continued fraction approximation will be closer than [294, Thrm. 171]. In the case where σ is of constant type, there is some maximum value for a_M and therefore a lower bound on the right-hand term, relative to $1/q_M^2$. I will return to this in the next chapter to explain why this undercuts our argument for unbounded growth.

From the above, we can write

$$\frac{1}{q_M^2 a_{M+1}}\left[1 - \frac{1}{a_{M+1}} + \mathcal{O}\left(\frac{1}{(a_{M+1})^2}\right)\right] \leq |\sigma_M - \sigma| \leq \frac{1}{q_M^2 a_{M+1}}\,. \tag{6.16}$$

Then, for genetic irrationals where a_{M+1} may be large and terms like $(a_{M+1})^{-2}$ and smaller neglected, we may write

$$|\sigma_M - \sigma| \overset{M\gg 1}{\approx} \frac{1}{q_M^2 a_{M+1}}\,. \tag{6.17}$$

6.3.2 Rational rotations

By *rotation* I mean a map on a circle of unit circumference that shifts all the points counter-clockwise by a distance σ. Beginning from a point on the circle x_0, after n steps the position will be

$$x_n(\sigma) = x_0 + n\sigma \mod 1\,, \tag{6.18}$$

with

$$\begin{aligned} n &= 0,\, 1,\, 2,\, 3,\, \ldots\,, \\ \text{and } 0 &\leq x_n < 1\,, \end{aligned}$$

In principle, $\sigma, x_0 \in \mathbb{R}$ but we will concentrate on $\sigma \in \mathbb{Q}$ for the moment and the whole part of x_0 has no bearing on the trajectory. This is a discreet flow, having both the properties (6.3a) and (6.3b) and is isomorphic to rotating the circle by angle $2\pi\sigma$.

First, let us consider a rational rotation, where $\sigma = p/q$ with $p \perp q$ and $p \in \mathbb{Z}, q in \mathbb{N}^+$, that is, they are coprime, whole numbers and q is positive. It is clear that the sequence is periodic, with a period q:

$$x_{n+q}(p/q) = x_n(p/q)\,, \tag{6.19}$$

7. It should be understood that the "goodness" of an approximation should always be considered relative to the length of q_M, called the *complexity* of the rational [294, p. 194] If this can get arbitrarily long, by choosing a higher M the approximation can always be made arbitrarily close.

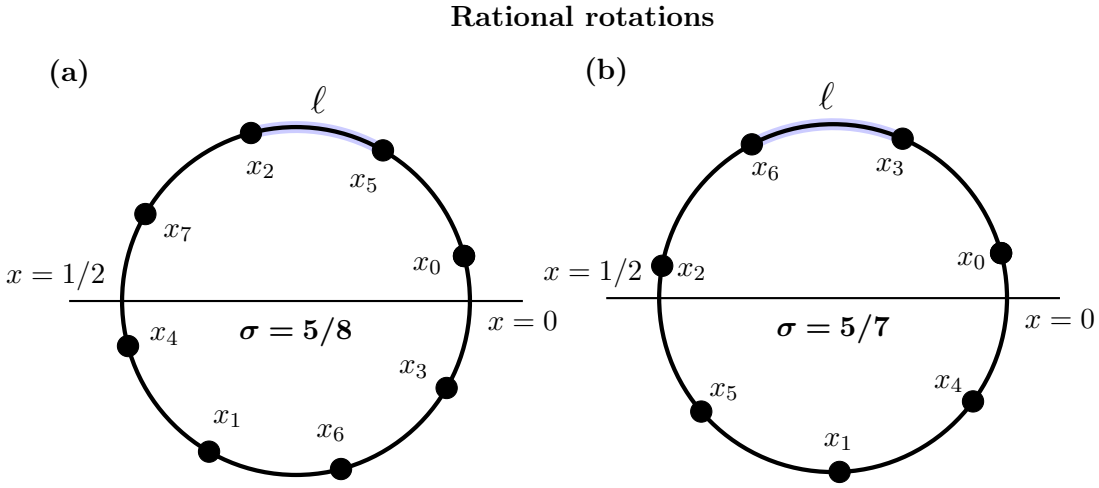

Figure 6.3 — An illustration for several statements about rational rotations made in the text. Both (a) (even q) and (b) (odd q) illustrate an assertion made in Section 6.3.2 that for a rational rotation of $\sigma = p/q$, the first q members of the $x_n(p/q)$ sequence lie on the circle in an equidistant manner, with the nearest-neighbor distance of $\ell = 1/q$, highlighted in blue. Subsequent points will fall on top of the first q.

since after a step of length $\Delta n = q$ the shift from x_0 will be the integer p and be sheared off by taking the fractional part. It follows immediately that any functions of x will also have the property (6.19).

For (6.18) with $\sigma = p/q$, q is also the shortest period. If there were a shorter period of say $m < q$, that would mean that $mp/q \in \mathbb{Z} \implies q \mid mp$, that q evenly divides mp.[8] But $\sigma = p/q$ is in its lowest terms by assumption and $m < q$, so must be that $q \nmid mp$ and there can be no orbit of length $m < q$.

Since (6.18) is a flow, this means the moment $x_n = x_{n'}$ for $n' \neq n$, the system falls into an *orbit*, of length q, and revisits all of the same points as before and in the same order. Thus, there can be no points revisited within an orbit because as soon as any point is revisited, the orbit has closed and no new points can be reached, according to equation (6.3b).

Since the orbit is of the same length whether $p = 1$ or not is easy to prove that the set of points in an orbit

$$\{x_0, x_1, \ldots, x_{q-1}\} = \left\{x_0, x_0 + \frac{1}{q}, x_0 + \frac{2}{q}, \ldots, x_0 + \frac{q-1}{q}\right\} \mod 1\,. \tag{6.20}$$

That is to say that the q points of the orbit divide the circle into segments of equal length,

$$\ell = \frac{1}{q}\,. \tag{6.21}$$

See Figure 6.3 for illustration.

What is more, we can say that when q is even, points in a rational rotation half an orbit apart will form antipodal pairs (i.e be separated by a length of $1/2$), for all values of p. The proof is similar to the one proving that q is the shortest orbit.

If q is even then there is an antipodal point $q/2$ steps away since

$$\frac{q}{2} \times \frac{p}{q} \equiv \frac{1}{2} \mod 1 \tag{6.22}$$

8. I use $m \mid (\nmid) n$ to mean m does (not) evenly divide n. Notations are in Appendix A.

because $p \in \mathbb{N}^+$. This can be the only point at that position with respect to where it occurs in the orbit because no points are repeated within an orbit.

When q is odd, on the other hand, a distance $1/2$ away from a point in the orbit is exactly between two other points. Since an orbit of length q divides the circle into q intervals of length $\ell = 1/q$, this is equivalent to saying that

$$\begin{aligned} \ell(m + \frac{1}{2}) &= \frac{1}{2} \quad \text{for } m \in \mathbb{Z} \\ \implies m + \frac{1}{2} &= \frac{q}{2} \end{aligned} \tag{6.23}$$

which is just to say again that q is odd.

6.3.3 Irrational rotations and how we will use them

A rotation such as (6.18) is *irrational* when $\sigma \in \mathbb{R} \setminus \mathbb{Q}$. This will never fall into orbits and has the property that

$$x_n \neq x_{n'} \quad \text{when} \quad n \neq n' \,. \tag{6.24}$$

The motion has been proven to be ergodic for any irrational σ (see [297] and the references at the start of Section 6.3), with an invariant measures being able to be defined on it. Take the Lebesgue measure of the real line interval $[0, 1[$ for an example of such a measure . Under the action of a rational rotation, a segment of length less than ℓ will fall into an orbit, just like any other set of points. The collections of segments it gets mapped onto over the course of an orbit will form an invariant set of positive, but not full measure. This is therefore not metrically transitive and therefore not ergodic. But in the irrational case, (6.24) ensures that orbits never close. That means that the set of points visited by a trajectory will be dense and a theorem by Weyl ensures that the points visited will be uniformly distributed when $n \to \infty$ [298, 299]. Thus, if one tries to form an invariant set by placing a segment centered around an initial point and including all the points visited in the trajectory of this segment, the invariant set is the entire circle, no matter how small the segment is. Therefore, no smaller invariant sets of positive measure can be made and irrational rotations are metrically transitive and so ergodic. Since this flow does not change the measure of the interval $[0, 1[$, we can say the flow is measure preserving and that $\mu(\mathrm{d}x) = \mathrm{d}x = \text{conts.}$.

Because irrational rotations are such simple ergodic systems, they are excellent for prototyping tools and arguments to apply to dynamical systems. The fact that rotations are related to reflections which in turn play such an important role on billiards is an argument for choosing them over some other simple, ergodic system like a harmonic oscillator. Our strategy is to construct series based on the observables in such a way that the ergodicity of the system and the behavior of those series imply things about each other. An example will make this clearer.

Consider a discrete, ergodic, observable, f, defined on a rotation such that

$$f(x) = \begin{cases} +1 & \text{for} \quad x \in \left[0, \frac{1}{2}\right[\\ -1 & \text{for} \quad x \in \left[\frac{1}{2}, 1\right[\end{cases} \,, \tag{6.25}$$

where x is as in equation (6.18). Because this is ergodic, it has the property that (6.5) converges to (6.2) and we have

$$\overline{f} = \langle f \rangle = 0 \,. \tag{6.26}$$

Then one can define a series, S_n, whose n^{th} term is $n\overline{f}_n$, which cannot grow faster than linearly. Taking the converse case where the ergodicity of f is in question, if one can define a series like S_n and then prove that it asymptotically grows more quickly than linearly,[9] then one has shown that the time average cannot converge to zero. If $\langle f \rangle = 0$ this would demonstrate that f is not ergodic. This is the essence of the idea we hope to employ with these counters and why we are so interested in their rates of growth.

Counters By *counter* we simply mean an object which counts some event in the underlying system, e.g. collisions with the billiard mirrors or only collisions of a certain type, etc. Some events may count negatively. Thus they are always integer valued and only ever change by ± 1. They are not themselves observables of the system as they are not just functions of the phase space. If one counts collisions with billiard mirrors for example, a given point in phase space could correspond to any value of the counter, depending on the history. As such, properties of the system like ergodicity or sensitivity to history are not directly inherited by counters and they must be linked to each other by arguments.

There are two kinds of arguments we can foresee making. One is after the fashion of the argument about S_n just above, where the counter would be constructed to behave like S_n (or its sum). The other would be an argument relying on metric transitivity and proving the counter is unbounded like the argument given above that irrational rotations are ergodic. The counter in that case could be the number of distinct angles visited. Metric transitivity and hence ergodicity then follows from the unboundedness of the counter by Weyl's theorem showing the points to be dense and equi-distributed.

6.4 Summary

In laying out the background theory for the next chapter I have taken several detours to draw connections to broader concerns so I will list here the key statements that must be kept in mind going forward. Terms are defined above near where they appear in *italics.*

- If a system is ergodic the infinite time average (6.4) converges to the ensemble mean (6.2) in the infinite time limit.
- Most physical systems evolve via a flow which has the properties (6.3a) and (6.3b).
- The measure of the phase space weights the ensemble average.
- Metric transitivity implies ergodicity and vice versa.
- If a point in phase space contains information about its trajectory for infinite time, the system is not ergodic.
- KAM theory says that the trajectories of integrable systems are confined to invariant tori in the phase space and this is generally incompatible with ergodicity.
- A billiard involves a virtual particle moving elastically inside of a boundary composed of mirrors. It can encode the dynamics of a one-dimensional system via coordinate transforms like those outlined in § 6.2.2.

9. i.e. that $S = \Omega(n)$, or to be explicit: $\limsup_{n\to\infty} \left| \frac{S_n}{n} \right| > 0$.

- Irrationals can be better and better approximated by series of fractions with denominators growing as in (6.14), leaving gaps from the rational approximated like (6.17).
- Irrationals can be generic, where the series of their continued fraction coefficients are unbounded or of constant/bounded type where this series has a maximum value.
- A rotation is defined by (6.18). Rational rotations fall into orbits and are therefore not ergodic while the opposite is true for irrational rotations.
- A counter is an integer valued object, distinct from a system's observables, that counts events that occur in the system, some events counting negatively.

Chapter 7

Parity counters

I will start this chapter by detailing the counter studied by Wang, Casati and Prosen in their 2014 paper [70] and the *parity counter* I worked on with M. Olshanii. Each is dealt with in their own section. Section 7.1 on the Wang-Casati-Prosen (WCP) counter is included for more than just background and explanatory purposes. I also offer some explanations for their numerical results and attempt to shed some light on what they mean for the underlying systems, both dynamical and physical. I argue that if their counter can be shown to undergo unbounded growth for an irrational right-triangular billiard, then those billiards are ergodic, and give reasons to suspect that that is the case.

With a slight modification, the WCP counter can be transformed into a counter which increments by the parity of the number of a certain kind of event in their system. This is what lead Prof. Oshanii to construct a different parity counter, defined over irrational rotations, as a way of trying to gain some insight. In section 7.2, I describe these counters (both the intermediate one and the one over irrational rotations), our estimation and numerical techniques, give our reasons for believing that it undergoes unbounded growth for generic type irrationals (defined in § 7.2.1), and treat its fractal dimension.

7.1 Wang-Casati-Prosen counter

The WCP counter, denoted K, is defined on a one-dimensional system of two hard-core point masses moving between impenetrable walls. When the particles collide with each other, K increases by one. When either particle collides with a wall, K changes sign. Let the positions of the masses m_1 and m_2 be x_1 and x_2 respectively, with their velocities then being v_1 and v_2. One can then make a coordinate transform $y_i \equiv \sqrt{m_i}x_i$, $w_i \equiv \sqrt{m_i}v_i$ where the index, $i = 1, 2$, just as in § 6.2.2.

As before, the particle masses determine the angle the billiard mirror which encodes the particle-particle collision makes with the coordinate axes. Let α be defined such that

$$\cos\alpha = \frac{m_1 - m_2}{m_1 + m_2}\,. \tag{7.1}$$

Then, if you allow θ to be defined such that

$$\tan\theta = \frac{w_2}{w_1}\,, \tag{7.2}$$

The Wang-Casati-Prosen billiard

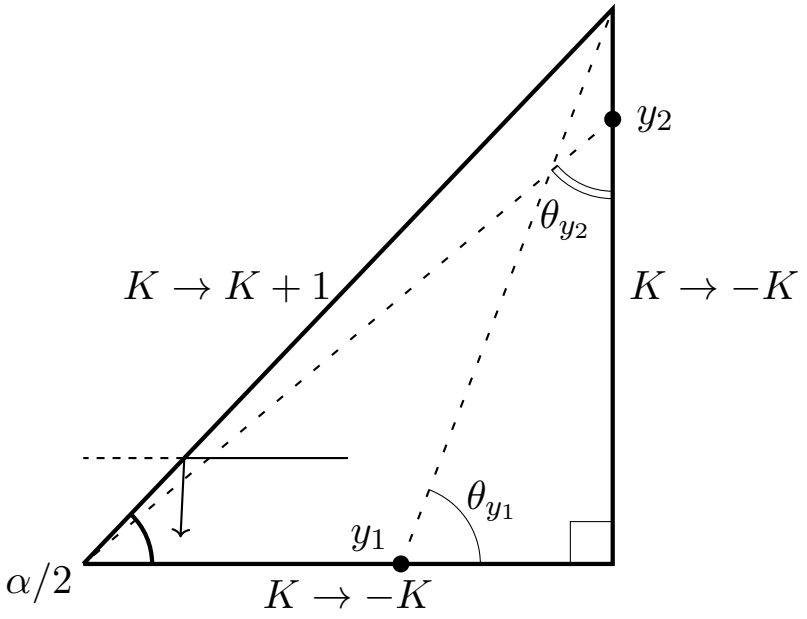

Figure 7.1 − The billiard on which the WCP is defined. Labels outside the sides of the triangle indicate how the counter changes when the virtual particle meets that side. The angles θ_{y_1} and θ_{y_2} are the angles subtended by the other leg when the particle meets a leg at y_1 and y_2 respectively. See 7.1.2 for the significance of this. The angle α is determined by the masses as in (7.1). Note that the trajectory colliding with the hypotenuse changes its direction by α.

after a particle-particle collision $\theta' = \alpha - \theta$. After a particle-wall collision, one the velocities, in both the transformed and one-dimensional space, will simply be inverted and $\theta' = -\theta$ or $\pi - \theta$.

Since a reflection changes the trajectory of a virtual particle by twice the angle of incidence, the billiard constructed from this system should feature:

i.) an angle of $\alpha/2$.

ii.) two mirrors on the coordinate axes to invert one velocity without changing the other.

iii.) only two dimensions, one for each particle.

This constrains us to the billiard shown in Figure 7.1: a right triangle with an angle of $\alpha/2$. It also lets us (or rather the original authors) define the WCP counter, K and imposes the rules by which it is incremented. The counter, K is defined so the velocity angle after some time, θ can be written

$$\theta = \varphi + K\alpha\,, \tag{7.3}$$

where $\varphi \in \{\theta_0, -\theta_0, \pi - \theta_0, \theta_0 - \pi\}$ and keeps track of what "quadrant" the velocity vector lies in. I say "quadrant" in quotes because they are not the four quarters of the circle delineated by the coordinate axes, but rather those made by two axes, one at angle α and the other at $\alpha+\pi/2$ (see Figure 7.2), about which the velocities will be distributed with reflectional symmetry. It absorbs changes in the velocity angle of $\pi/2$ leaving the changes by α to be accounted for.

Quartering of the circle of the velocity angle by φ

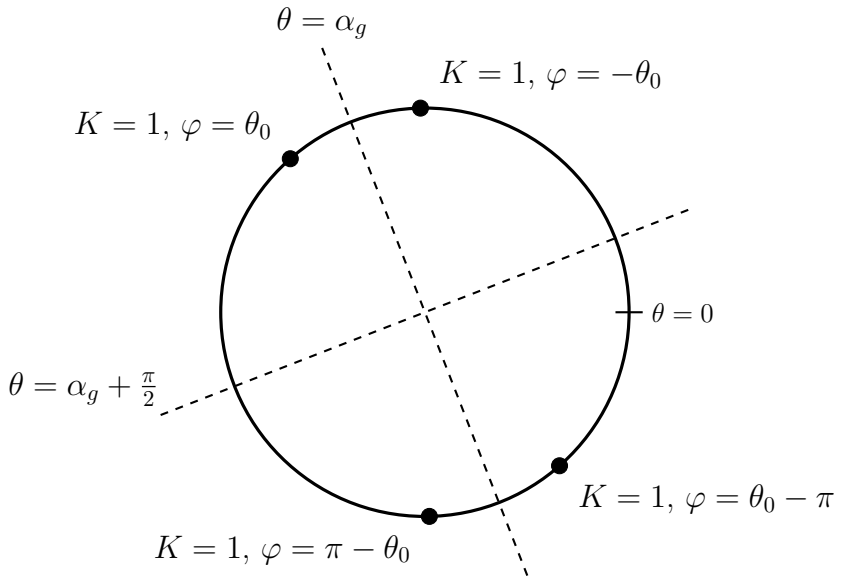

Figure 7.2 — Diagram illustrating the function of φ in relating the WCP counter K to the velocity angle, θ, in equation (7.3). Here the "strongly irrational" of $\alpha = \alpha_g = \phi\pi$ case is illustrated. This figure also makes clear that as $K \to \infty$, the velocities will be distributed in four sets, related to each other by reflection about the two axes indicated by the dashed lines.

You will note that apart from this, it resembles a rotation, similar to equation (6.18), but here K is incremented in a special way.

The rules for changing K to K' after a collision are written

$$\begin{aligned} K' &= K + 1 \quad \text{(hypotenuse/particle collision)} \\ K' &= -K \quad \text{(leg/wall collision.)}\,. \end{aligned} \tag{7.4}$$

When α is rational multiple of 2π, K has a certain set of values it can visit because it captures the transformation of the velocity angle by rotation of 2α (see Figure 7.1). As I said in § 6.3.2, when $\alpha/2\pi = p/q$ for positive integers p, q, there are only q angles that can be reached. If the initial angle θ_0 is irrational, the trajectory itself may not fall into an orbit [293, §4.1] but there is a limited number of directions the velocity can take, n_K^* for a rational $\alpha/2\pi = p/q$ that is an invariant for a particular q. This does not impose a maximum value on K however since all values of K are equivalent modulus the orbit length, q. This is not true when α is irrational however. In that case, every value of K corresponds with a unique value of θ, up to rotations of $\pi/2$, absorbed by φ. To see this let

$$\tilde{\theta} \equiv K\alpha \mod 2\pi\,. \tag{7.5}$$

The only way for $\tilde{\theta} = K'\alpha = K\alpha$ for $K' \neq K$ is for there to be some non-zero integer $q : K\alpha \equiv (q+K)\alpha \mod 2\pi \implies q\alpha \equiv 0 \mod 2\pi$. But this implies that $\alpha = p/q \cdot 2\pi \in \mathbb{Q}2\pi$ for $q \in \mathbb{N}^+$, $p \in \mathbb{Z}$, in contradiction with $\alpha/2\pi \in \mathbb{R} \setminus \mathbb{Q}$. For φ to cause a collision like the one mentioned above it must hold that

$$(K - K')\alpha \equiv \frac{p}{2}\pi \mod 2\pi, \tag{7.6}$$

where $p = 0, 1, 2$ or 3. But that is to say that an integer times an irrational is rational, an impossibility.

The main result of the paper [70] is numerical evidence about what happens when, α is "strongly" irrational, meaning that it is a number far from rational times π. As an example they take the golden ratio,

$$\phi = \frac{\sqrt{5}-1}{2} \tag{7.7}$$

and set $\alpha_g/\pi = \phi$. They then study approximants to this, $\alpha/2\pi = p/q$ in a manner I will discuss in § 7.1.1. They find numerically that, in the strongly irrational case where $\alpha \to \alpha_g$, while K has a nonzero probability of reaching any value on the number line, the probability that it reaches any particular number falls off exponentially as one moves away from the initial value, K_0. Since the average of this over time, $\bar{K}$, will just be the initial value, K_0, the velocity angle, θ, must also be localized in time around θ_0, at least up to rotation by $\pi/2$. But the average over K_0 will be ill-defined, so when α is "strongly" irrational the billiard is not ergodic. They claim that this is a novel kind of non-ergodicity, distinct from the kind predicted by KAM theory where the system is restricted to tori of positive with associated, distinct invariants. This is what keeps the system from being metrically transitive and therefore makes it non-ergodic.

The "weakly" irrational case is more strictly defined as being those billiards proven to be ergodic by Vorobets in 1997 [274]. In his proof, Vorobets constructs "weakly" ergodic billiards as those billiards where the irrational angle α and the approximant, $\alpha_q/\pi = p/q$, differ by less than

$$\left(2^{2^{2^q}}\right)^{-1}. \tag{7.8}$$

This particular rate of approximation is particular of Vorobets' proof and was not chosen for its general significance. "Strongly" irrational in this context means that the difference is greater than (7.8).

As an example they take $\alpha/\pi = L$ where

$$L \equiv \sum_{k=1}^{\infty} 10^{-k!}, \tag{7.9}$$

Liouville's constant, an irrational number very close to a rational one. The number L is a generic irrational and we can find in its continued fraction representation a coefficient, $a_{v+1} = 10^{(v+1)!-2v!}, 2 \mid v$ [300], which can be arbitrarily large and so the error (6.15) can be made arbitrarily small. Wang et al. find that billiards with this angle behave ergodically, in line with [274].

In the case of near-rationals like $\alpha/\pi = L$, Wang et al. find that all of the velocity angles allowed by a rational approximant are visited before the approximate sequence diverges from the true one. They say this is a requirement of ergodicity and the fact that this does not happen in the strongly irrational case is evidence of non-ergodicity and suspect that the cross-over from weakly irrational and ergodic and strongly irrational and non-ergodic is just when the approximation begins to fail before all of the allowed angles are visited. This is similar in feel to our reason for insisting on a generic irrational in our system: we need to ensure that certain cycles in the approximate sequence generated by a rational have been completed before the next rational convergent is used.

7.1.1 Note on numerical methods

It is impossible to exactly, indefinitely simulate a system which depends on an irrational number with a machine operating at finite precision. I will spend this section comparing their methods to our own as we use the same trick for determining the range of validity for our calculations. It will also give me chance to discuss a few shortcuts we take that allow our simulations to be more efficient, since there are properties they were interested in that we are not.

Both investigations study systems that can be modeled discretely. They keep track of which side of the triangle the virtual particle reflects off of and construct a code for a trajectory consisting of the sequence of these reflections. We work with a rotation (6.18), which is explicitly discrete. They also simulate the whole trajectory through phase space as exactly as possible and derive discrete sequences from that. We have a phase space of smaller dimension and need only the discrete series (phase points, counter values, etc.) associated with the rotation.

A finite-precision calculation must work with rational numbers and so can only approximate a system based on an irrational number. For trajectories like the ones under consideration in both our systems, that means a rational approximation will give the same string of outcomes as the "true" sequence based on the irrational number until the difference between the rational approximant and the irrational approximated becomes significant enough to affect the sequence of discretized outcomes and the sequences differ. Since we do not have access to this exact sequence we need something else to compare to. Since methods for constructing a sequence of rational numbers that are better and better approximations of a given irrational are old and reliable (continued fractions being a prime example. See [294, §10.15] and references therein), we can construct a series of rational angles that yield sequences of outcomes closer and closer to the true sequence based on the irrational of interest when fed to our simulation. To approximate a given irrational, they generate a series of approximants, p_n/q_n indexed by n. The angle used for the billiard is then the rational angle $\alpha/2\pi = p_n/q_n$. An approximant based on α_q will give a sequence that differs from one based on $\alpha_{q'}$, at time $\tau^*(q, q')$ where $q' > q$. At that time , we can no loner trust the sequence generated by α_q since approximants with larger denominators approximate the sequence better and better. Therefore, if we have two rational approximants, the sequence of outcomes they give can be relied on to mimic the true sequence insofar as they agree with each other. After they diverge from one another, nothing can be said with certainty until the sequence from a better rational approximant is generated.

Both Wang at al. and my co-authors and I use this method for generating sequences of counter values which are exact up to some time step. The way we generate these sequences based on approximants is a little different however. They were interested in the dynamics of the billiard underlying their counter. They went so far as to randomize the position, but not the velocity angle, of the virtual particle after each reflection to gain some insight into how the particular geometry of the billiard affected the counter sequence [257]. This means they must run a two-dimensional billiard simulation which keeps all of the position information intact. Our system is simply a rotation, so we do not need to simulate the circle, radius and all, in order to get our results. We simply use a software package that can work in exact rationals (either Mathematica or the "fractions" library in Python) and solve an equation which reads like equation (6.18) over and over again, calculating our counter based on the resulting phase point. Because our underlying system has very little dynamics apart from those that affect the counter, our calculations can be relatively simple.

7.1.2 Explanation of the results

Up to now I have not said anything about the work of Wang et al. that they have not published themselves [70, 257], only offered it as background inspiration and used it to introduce a few of the features of our own work on counters. In this section I will offer my own explanations of some of their results. First I will attempt to shed some light on the behavior of one of their observables, n_K^*, the number of velocity angles visited by the time their approximation fails, τ^*, since it will be a good way for me to introduce a few remarks about the localization of their counter, K.

They calculate how this n_K^* value grows with q_n, that is with better and better approximations of $\alpha_g/\pi = \phi$. What they are after here is how the number of velocity angles visited by the trajectory of the truly irrational billiard grows with time. They find that $n_K^* \sim \ln(N)$, consistent with their main finding that K is localized by an exponential around K_0.

Asking how n_K^* grows with q_n is equivalent to asking how the maximum K value, $K_{\max}$, grows with time for $\alpha = \alpha_g$ since the number of distinct velocity angles visited at any time is $4K_{\max}$, since φ can take on up to four values. The question of how $K_{\max}$ increases is a question of the rate of leg hits coming one after the other. Imagine K is at some positive value less than $K_{\max}$. After a $K \to K+1$ hypotenuse collision, the virtual particle must strike a leg and $K \to -K$. It can then either strike the hypotenuse again, increasing by one and decreasing its absolute value, or it may hit the other leg taking it back to its previous positive value. The next collision must necessarily be with the hypotenuse and increase the absolute value. These double-leg hits are the only way $|K|$ can grow as if this never happened it would go: $K = 0 \to 1 \to -1 \to 0$ forever. This would correspond to $\alpha = \pi/4$ and θ_0 being on the horizontal or vertical axis. It has already been shown that almost all trajectories starting perpendicular to one of the legs of the triangle form closed orbits [275]. This is avoided by Wang et al. even in the rational case by insisting $\theta_0/\pi \in \mathbb{R} \setminus \mathbb{Q}$. However, for n_K^* to grow, the leg-leg hits (let this be denoted 'LL' with a hypotenuse hit putting an 'H' in the sequence) must occur when $K = K_{\max}$, so it is little wonder that n_K^* grows slowly.

Looking more carefully, we can say more about the rate than that it is "slow" and see how strongly their results suggest non-ergodicity. After hitting the hypotenuse and then a leg, let the probability that the particle will hit the other leg be p. If it does hit the other leg, then it then must hit the hypotenuse and then a leg (giving a sequence of $HLLHL\ldots$), incrementing the value of K by one and returning us to a similar situation: the virtual particle against one leg with a probability, p', of hitting the other leg. We can imagine averaging over a great many values of p to arrive at a $\bar{p}$ which, if the billiard is ergodic and we average long enough, will converge to a definite value, $\langle p \rangle$, where the ensemble average is taken over all points on the legs. We can then say that, on average, the probability of getting two leg-leg hits in a row (albeit with the required hypotenuse hit in between) is then $\bar{p}^2$. The average probability of getting a string of hits like $HLLHLLHL\ldots$, with the K-incrementing sequence, 'LLH', repeated k times being $\bar{p}^k$.

If we let $K_0 = 0$, then the time average of K, $\bar{K} = 0$, and it will take, on average, $K_{\max}+1$ sequential leg-leg hits to raise the value of $K_{\max}$. We can say that the probability of getting k double leg hits in a row in time t is then $p_k(t) = \bar{p}^k t$. If one fixes k (and $K_{\max}$ will be fixed until this string of $k > K_{\max}$ hits occurs) and takes the long-time limit, $p_k \to 1$. Now since the meaning of "$\sim$" is "ratio approaches one" we can say that, in the long time limit, $\bar{p}^{-k} \sim t$. With some rearranging and by recalling that $t \sim \tau^* \sim N$ [70, Fig. 2] and that $k \sim K_{\max} \sim n_K^*$ we may write in their notation

$$n_K^* \sim \log_{\frac{1}{\bar{p}}} N\,. \tag{7.10}$$

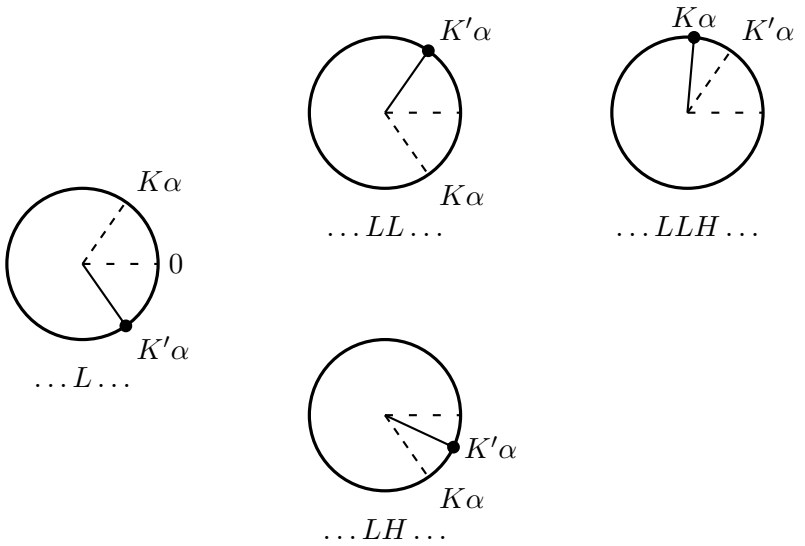

Figure 7.3 — Illustration of the statement that $|K|$ only grows on LL hits. On the left, the billiard undergoes a leg hit, L and $K \to -K$ (7.4). Next, it either hits the hypotenuse (LH) and $|K|$ goes down by one *or* it hits the other leg first (LLH) and $|K|$ goes up by one. This is the connection to parity counters captured by equation (7.14).

This explains why n_K^* grows logarithmically which is equivalent to saying that K is localized around K_0 by an exponential. It also means that their numerical results suggest that $1/\bar{p} \approx e$.

To try and get a sense of how much the probability, p, is due to the geometry of the billiard, let us try and approximate $\langle p \rangle$. Let us assume that the particle is equally likely to hit any spot on the leg, and is equally likely to go in any direction, which is to say that we assume that the underlying billiard is at least quasi-ergodic. Wang et al. do not deny this. In fact, the novelty of their discovery is a system that seems quasi-ergodic, and yet the averages do not converge. Under this assumption the probability that the particle will hit one leg after the other will be equal to the angle subtended by the other leg from a point on the leg it just hit divided by π, the total angle it will scatter into. Let the hypotenuse be of unit length and let θ_{y_2} be the angle subtended by the other leg when the particle is on the leg opposite the angle equal to $\alpha/2$ and vice-versa for θ_{y_1} (see figure 7.1 for illustration and definitions.) Some geometry then gives

$$\theta_{y_1}(y_1) = \tan^{-1}\left(\frac{\sin(\alpha/2)}{\cos(\alpha/2) - y_1}\right), \quad y_1 \in [0, \cos(\alpha/2)] \tag{7.11}$$

$$\theta_{y_2}(y_2) = \tan^{-1}\left(\frac{\cos(\alpha/2)}{\sin(\alpha/2) - y_2}\right), \quad y_2 \in [0, \sin(\alpha/2)]\,. \tag{7.12}$$

Averaging together the average values of (7.11) and (7.12) gives an estimate of $\langle p \rangle$ written

$$\langle p \rangle = \frac{\langle \theta_{y_1} \rangle_{y_1} + \langle \theta_{y_2} \rangle_{y_2}}{2\pi}\,. \tag{7.13}$$

Calculating the reciprocal of this numerically gives $\bar{p}^{-1} = \langle p \rangle^{-1} = 2.9876\ldots$. This differs from e by only 10% and gives a growth rate for n_K^* which differs from a natural log growth rate by a factor of $\ln(\bar{p}^{-1}) = 1.0945\ldots$, which is not easily distinguished.

This hypothesis is supported by the numerical evidence that they published in [257]. They ran a modified simulation where after every collision, they randomized the position of the virtual particle in the billiard, exactly removing the assumption behind the above calculation. They found that exponential localization disappeared [257, Fig. 5] and the distribution of the probability of K at a given time step, $P(K,t)$, followed a Gaussian distribution with a variance $\langle K^2 \rangle = Dt$ with $D = 1.04$ at $t = 10^7$. For comparison, $\langle K^2 \rangle$ for the unmodified system at $t = 10^7$ was on the order of 10^2, a difference of five orders of magnitude. So it is possible that the localization of the counter is explained by the geometry of an underlying, ergodic billiard.

Wang et al. rightly observe that a trajectory can be dense in the phase space but not ergodic. For example it could be like the Cantor set, $\mathcal{C}$, which is dense on the interval $[0,1]$, but has gaps. That is to say that if one places sub-intervals of length ε centered on each $c \in \mathcal{C}$ there will be values of $\varepsilon > 0$ such that subsets of the interval $[0,1]$ with positive measure will not be covered by the sub-intervals centered on the elements of $\mathcal{C}$, even though there is no positive value of ε so small that the sub-intervals will not contain at least two members of $\mathcal{C}$; because $\mathcal{C}$ is dense. If the trajectory took points in $\mathcal{C}$ to each other, one could make an invariant set by taking the union of the points in these sub-intervals with $\varepsilon > 0$ and the system would not be ergodic. However, nothing like this can be the case for K.

Let $\{\tilde{\theta}\}$ be the set of angles in the quadrants quartered by φ visited by a trajectory. If α is irrational then equation (7.5) is just an irrational rotation and by Weyl's theorem [298], the angles in the set $\{\tilde{\theta}\}$ are not just dense when $K \to \infty$ but also equi-distributed because this set will be homomorphic to the ones visited by a typical irrational rotation. Therefore there are no gaps of the kind described above in the Cantor set and it, as already shown [261, Prop. 2.16], is ergodic. Therefore, since each value of K corresponds to a unique angle in the set as argued above, to show that θ is not ergodic (and if it is not, the billiard cannot be), they must show that K is bounded from above and so can be used to define an invariant set *or* make some argument involving φ, which they do not. They say K is not so bounded and as long as $\bar{p}$ in (7.10) does not decrease faster than N^{-1}, it cannot be. Therefore, the only invariant set under (7.5) is the continuous angle measuring $\pi/2$, which has full measure, meaning the system is metrically transitive and therefore ergodic by Birkhoff's theorem [64].

Wang et al. also claim that the non-ergodicity which they deduce from the localization of K is of a new type, distinct from the type predicted by KAM theory. KAM theory states that a closed system is either ergodic, or its phase space is divided into toriodal regions with differing invariants. Here they have a phase space which is undivided, all values of K, and therefore θ, are accessible from every initial condition, and the phase space is densely covered, yet $\bar{\theta} = \theta_0 \neq \langle \theta \rangle$, so it is not ergodic. But their evidence for this last part are finite numerical simulations. Ergodicity only requires the average to converge after infinite time; this is why metric transitivity depends on whether the sets are strictly invariant under the flow. If K is unbounded for all θ_0, as they suspect but do not show, then the infinite number of velocity angles visited during a trajectory are dense and equi-distributed by Weyl's theorem [298]. This means any interval of angles, no matter how small, will cover every part of the circle eventually under the action of the flow. This means the only invariant set is all the angles in the circle so the flow is metrically transitive and the billiard is ergodic, at least in the velocity angle. Since there is no proof either way, only numerical evidence, the question remains open.

7.2 Product or parity counter

The counter detailed in the last section, 7.1, served as inspiration for a counter constructed by Prof. M. Olshanii. The hope is to use the techniques developed here to analyze other parity counters defined over other dynamical systems. For instance we hope to argue from the growth of our counter to the ergodicity of strongly irrational right triangular billiards via an intermediate counter to the WCP counter. In this section, I will describe the counter, connect its growth to the behavior of its underlying system, an irrational rotation, and give numerical evidence that its growth is unbounded when the irrational is generic.[1] I will also discuss the fractal dimension of its graph.

Calling this new counter a "parity" counter is the clue to its connection to the WCP counter. To that end I introduce the modified WCP counter which serves as a bridge between that one and our parity counter. Let this intermediate counter be κ, defined by rewriting equation (7.3) to read

$$\theta = \varphi + (-1)^n \kappa\alpha\,, \tag{7.14}$$

where the index $n = 0, 1, 2, 3, \ldots$ counts the number of collisions, beginning with the first particle-particle (i.e. hypotenuse) collision. That is to say that $n = 0$ until the first particle-particle collision, whereupon in changes to one, and increments by $+1$ on every collision, of whatever kind, thereafter. After this change to the counter, κ remains unchanged after particle-wall (leg) collisions. However, after particle-particle collisions, it increments by ± 1 depending on the parity of number of times both particles undergo particle-wall collisions before colliding with each other. In billiard language: on hypotenuse hits, κ increments by the *parity* of the total number of leg hits. This has the same relation to θ as K does in terms of its ergodic properties, localization, etc.; it just counts slightly differently.

The counter I treat here also counts the parity of the total number of events of a certain kind in an underlying system. However, the system is not a billiard but an irrational rotation as discussed in Section 6.3. Irrational rotations are a fairly simple dynamical system and known to be ergodic, so they provide a way to explore the properties of parity counters and find methods that can be applied more generally. I will define these objects in the next section.

7.2.1 Notation and mathematical grounding

First I will go over the specifics of the notation I will use, collect the relevant mathematical facts from Section 6.3, and demonstrate a few results that I will be relying on in the rest of this section. For the convenience of having everything in one place, I will lay out all of the relevant quantities here, even if they have already been defined in Chapter 6.

Let the position of an irrational rotation trajectory be denoted x, which falls in the real interval $[0, 1[$. Then it evolves via a discrete flow so that at time n (that is, after n steps) it takes the value

$$x_n(\sigma, x_0) \equiv x_0 + n\sigma \mod 1, \tag{7.15}$$

with

$$\begin{aligned} n &= 0,\, 1,\, 2,\, 3,\, \ldots\,, \\ \sigma &\in \mathbb{R} \setminus \mathbb{Q}\,, \\ \text{and } x_0 &\in \mathbb{R}\,. \end{aligned}$$

1. Defined in (6.10) or 7.22 in this section.

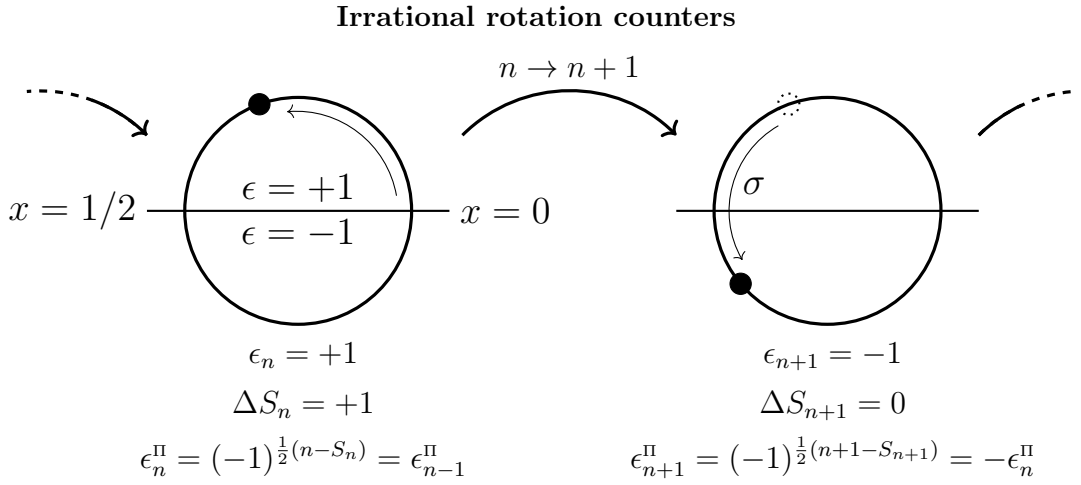

Figure 7.4 − A cartoon of an irrational rotation and two counters ((7.35) and (7.18)) defined over it. Notice that ΔS^{Π}_n (which is reckoned relative to the step just prior to the two pictured) cannot be so easily written down since every ϵ^{Π} depends on the one before it, the change in the sum between any two vales of n depends on the entire history of the trajectory, whereas ΔS_n does not. See text for a fuller explanation.

Again, the whole part of x_0 is irrelevant since everything is taken modulus one. This construction makes x_n a real number between zero and one. Let the invariant measure over this space, $\mu(x)$, be the Lebesgue measure.

Let the increment of a counter be, ϵ_n and behave as, f_n (6.25). Let that counter be S_n, so that

$$S_n(\sigma, x_0) = \sum_{n'=1}^{n-1} \epsilon_{n'}(\sigma, x_0) \tag{7.16}$$

$$\text{where} \quad \epsilon_n(\sigma, x_0) = \begin{cases} +1 & \text{for} \quad x_n \in \left[0, \frac{1}{2}\right[\\ -1 & \text{for} \quad x_n \in \left[\frac{1}{2}, 1\right[. \end{cases} \tag{7.17}$$

The parity counter is then the sum of product over the steps of this counter. This reads

$$S^{\Pi}_n(\sigma, x_0) = \sum_{n'=1}^{n-1} \epsilon^{\Pi}_{n'}(\sigma, x_0) \tag{7.18}$$

$$\text{where} \quad \epsilon^{\Pi}_n(\sigma, x_0) = \prod_{n'=0}^{n} \epsilon_{n'}(\sigma, x_0) \tag{7.19}$$

$$= (-1)^{\mathcal{N}^-} , \tag{7.20}$$

and $\epsilon_n(\sigma, x_0)$ is as defined in (7.17). The integer $\mathcal{N}^-$ is the number of instances of $\epsilon_{n'}(\sigma, x_0) = -1$ in the interval $0 \leq n' \leq n$. This means $\epsilon^{\Pi}_n(\sigma, x_0)$ depends on the parity of instances of $\epsilon_{n'}(\sigma, x_0) = -1$ up to and including the time, n: it is $+1$ for a even number of instances and -1 for an odd number. This has the effect of giving $S^{\Pi}_n(\sigma, x_0)$ a memory in the sense discussed in section 6.1. Each value of $\epsilon^{\Pi}_n(\sigma, x_0)$ depends on the value just before it and so cannot be reconstructed without it, and the same for that one, and the one before that, and so on back

to the first. To calculate the value of $S_n^{\Pi}(\sigma, x_0)$ at a given n, one needs a complete series of either $\epsilon(\sigma, x_0)$ or $\epsilon^{\Pi}(\sigma, x_0)$.

A cartoon illustrating an irrational rotation, $\epsilon_n(\sigma, x_0)$, the change in its sum, denoted ΔS_n, and the counter in 7.18, $S_n^{\Pi}(\sigma)$, is presented in the Figure 7.4, with detailed explanations to follow.

Similar to Wang et al in [70], discussed in the previous section, we will be studying the behavior of our counter, (7.18), by constructing a series of rational approximants and seeing how the counter behaves as these become more precise. The continued fraction expansion of the shift σ is

$$\sigma = [a_1;\, a_2,\, a_3,\, a_4,\, \ldots] = a_1 + \cfrac{1}{a_2 + \cfrac{1}{a_3 + \cfrac{1}{a_4 + \ldots}}}\,, \tag{7.21}$$

with $a_1 \in \mathbb{Z}$ and $a_i \in \mathbb{N}^+$ for $i > 1$. Recall a *generic irrational* is one where the corresponding sequence of continued fraction coefficients, a_1, a_2, a_3, ..., is not bounded from above, i.e.

$$\text{generic irrational} \implies \forall A \in \mathbb{N} \quad \exists\, i : a_i \geq A. \tag{7.22}$$

If, on the other hand, the corresponding sequence, a_1, a_2, a_3, ..., *is* bounded from above, then it is a *constant type irrational*, i.e.

$$\text{constant type irrational} \implies \exists\, A \in \mathbb{N} : \forall i\ a_i \leq A. \tag{7.23}$$

The methods that follow only hold for generic irrationals where the sequence of coefficients is unbounded.

Conceptually, the difference between them is that generic type irrationals can be approximated arbitrarily well, taking into account the magnitude of their denominator, by their continued fraction while constant type irrationals cannot. The magnitude of a continued fraction coefficient can be thought of as indicating the precision of the expansion at that point [294, §11.8]. Since the coefficients of a constant type irrational are bounded from above, it can only ever be approximated by its continued fraction finitely well. The archetypal example being the basis for the "strongly irrational" case discussed in the last section, $\phi = [0; 1, 1, 1, \ldots]$. It is has the strictest bound, where A from (7.23) takes the value one and its distance from the closest rational is the greatest possible [301, Thrm 1.5].

Since we are going to be studying $S^{\Pi}(\sigma)$ by studying rational convergents of an irrational σ, let a rational approximant of a given σ be σ_M. The approximant is formed by truncating the sequence of continued fraction coefficients, $a_1, a_2, a_3, \ldots, a_i, \ldots$, at $i = M$. This gives

$$\sigma_M = [a_1; a_2,\, a_3,\, \ldots,\, a_M] = a_1 + \cfrac{1}{a_2 + \cfrac{1}{a_3 + \cfrac{1}{\ddots + \frac{1}{a_M}}}} = \frac{p_M}{q_M},$$

with p_M and q_M nonzero, coprime integers.

The study of how continued fractions are constructed and how the sequence of their approximants converges is it's own domain [302] and using it to find interesting examples

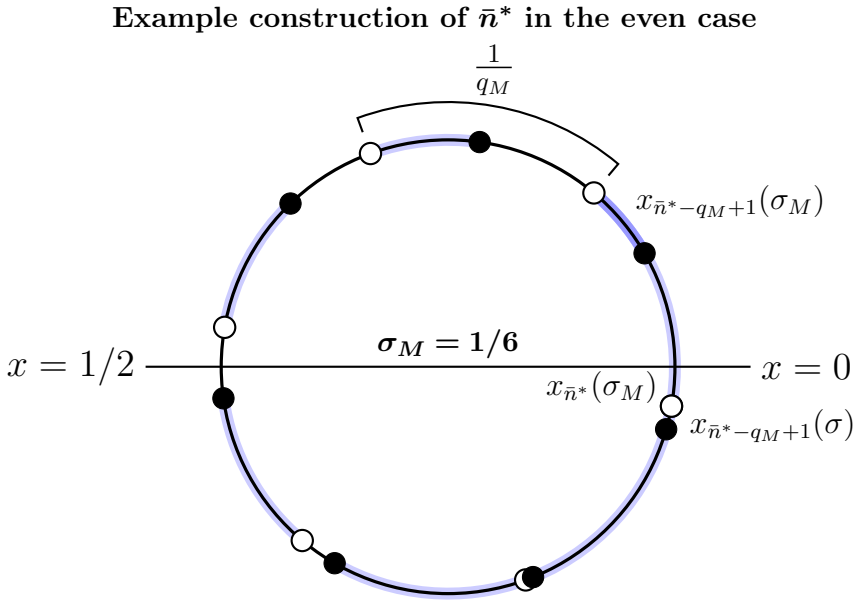

Figure 7.5 — An illustration of $\bar{n}^*$ for $2 \mid q_M$ and $\sigma < \sigma_M$. The size of $|\sigma - \sigma_M|$ is exaggerated for illustration. Points generated by the approximant rational rotation are in white while the points from the irrational rotation are in black. The difference between points generated by the same n is highlighted in blue, reckoned always counterclockwise. The point at $x_{\bar{n}^*}$ is not pictured. Here, since points generated by the rational rotation are antipodal, to ensure condition (7.27) for all values of x_0, some gap between a white point and a black one must cover the whole length of $1/q_M$.

of irrational and their approximants may be interesting but we take a more straightforward approach. Since our proof relies on σ being a generic irrational, we will just construct a generic irrational to do our numerics with. We simply take $a_i = i^2$ giving

$$\sigma = [1;\, 4,\, 9,\, \ldots,\, i^2,\, \ldots] \approx 1.24328847839971564\ldots\,. \tag{7.24}$$

This lets us say something more about the growth of the denominators of the approximants, q_M, grow with M than equation (6.14) in the last chapter, which was true for all truncated continued fractions. We can see that for the particular generic irrational we studied

$$\frac{q_M}{q_{M-1}} = M^2 + \frac{q_{M-2}}{q_{M-1}} \overset{M\gg 1}{\approx} M^2\,. \tag{7.25}$$

This has as a corollary,

$$q_M \overset{M\gg 1}{\approx} \text{const} \times (M!)^2\,. \tag{7.26}$$

7.2.2 Error introduced by rational approximation

Two approximant sequences, $S_n^{\Pi}(\sigma_M)$ and $S_n^{\Pi}(\sigma_{M+1})$, will match up to some n. Up to this point, we can also be sure that they both match the "true" sequence, $S_n^{\Pi}(\sigma)$. This is the same method used by Wang et al. to determine their time of failure, which they call τ^* and we call n_M^*, and is the how we could be certain of our calculations of $S_n^{\Pi}(\sigma)$ up to various

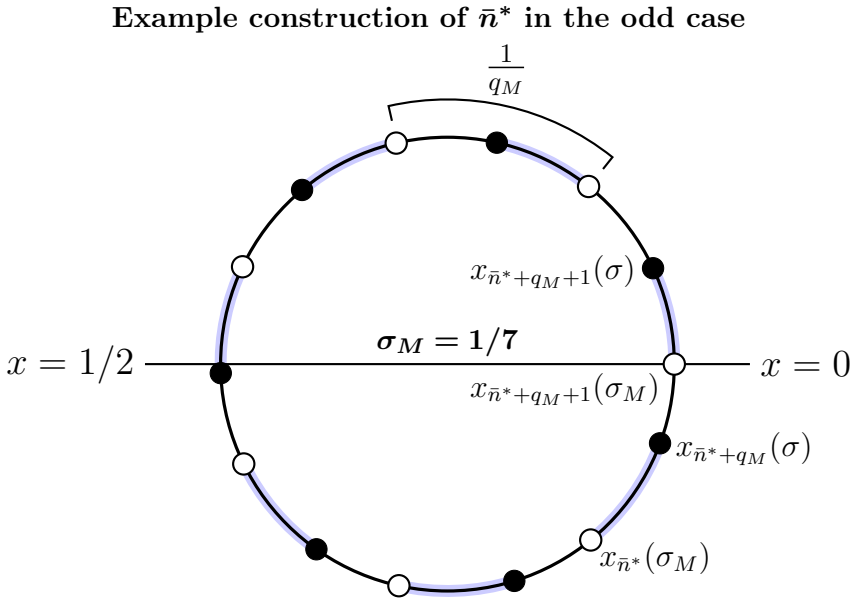

Figure 7.6 — Similar to Figure 7.5 but for $2 \nmid q_M$ and $\sigma > \sigma_M$. As before, rational rotation points are in white and irrational rotation ones are in black, difference between points generated by the same n in blue. Here, points generated by the rational rotation are opposite a point halfway between two other points, so to ensure condition (7.27) for all values of x_0, some gap between a white point and a black one needs only cover a length of $1/2q_M$ to ensure no diameter can be drawn without falling into a blue gap.

times. Since this will be longer with larger M, it lets one consider M as a degree of precision for an approximate sequence; a less "precise" sequence being one that fails at lower n. This relationship can be calculated exactly up to whatever n, M one has the computational and temporal resources for, but an estimate of this value will be useful. If you can be sure that you do not go over n^*, you can calculate a single sequence and be twice as efficient. It also will provide some insight into the long-term behavior of $S_n^{\Pi}(\sigma)$.

To see how the precision in terms of n grows with M, let us define $n_M^*(\sigma)$ as the smallest n such that

$$\epsilon_n(\sigma_M, x_0) \neq \epsilon_n(\sigma, x_0). \tag{7.27}$$

However this has two undesirable features. First, it must be calculated by calculating the series $\epsilon_n(\sigma, x_0)$ and $\epsilon_n(\sigma_M, x_0)$ explicitly, starting from $n = 1$, to compare the terms until the two that do not match are found. Second, it depends on the initial condition, x_0. So I will construct an estimate of n^*, which I will call $\bar{n}^*$, which is both simple to calculate and independent of x_0. Sadly, this will over-estimate n^* since it must hold for all initial conditions. I do this for the case of *generic* irrationals since those are the ones our model for the growth holds true for.

Let $\bar{n}_M^*$ be the lowest n such that the error, $|\sigma_M - \sigma|(\bar{n}_M^* + q_M)$, is large enough such that there is no diameter of the circle that does not fall in between $x_n(\sigma_M)$ and $x_n(\sigma)$ for some $n \in [\bar{n}_M^*, \bar{n}_M^* + q_M - 1]$. This guarantees that over the course of the next orbit of $x(\sigma_M)$, starting at $\bar{n}_M^*$, $\epsilon(\sigma_M)$ must differ from $\epsilon(\sigma)$ for any initial condition. Looking ahead an orbit should help to mitigate the over-estimation that comes from the requirement that it work for all initial conditions. The length that must be covered will be a different depending on

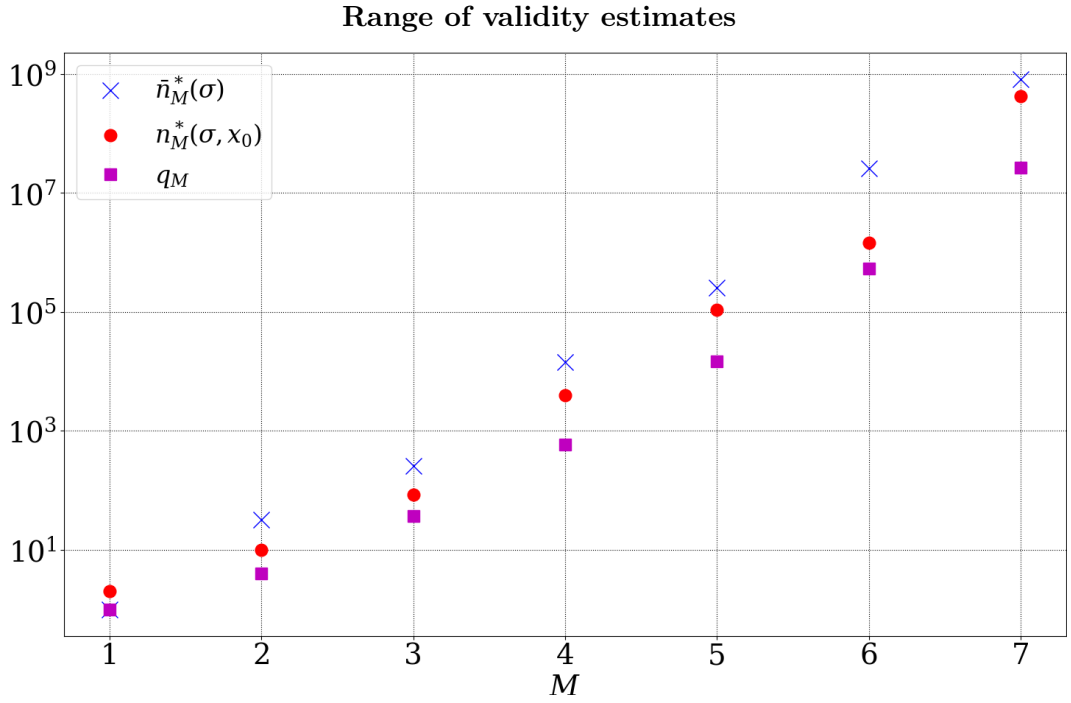

Figure 7.7 — The first time, $n^*_M(\sigma, x_0)$, that a convergent σ_M for the rotational step σ gives a value of $\epsilon_n(\sigma_M)$ that is different form the one predicted by its exact value, $\epsilon_n(\sigma)$ is given in red dots. An estimate, $\bar{n}^*_M(\sigma)$ (defined in equation (7.31)), is shown in blue crosses. The orbit length of the approximant, q_M, is given in magenta squares. Points in red are given by (arithmetically exact) numerics. It is crucial to our method, given in § 7.2.4 that n^*_M always be greater than q_M, hence our insistence that σ be a generic irrational. See that section for details.

whether q_M is even or odd. When q_M is even as in Figure 7.5, $|\sigma_M - \sigma|(\bar{n}^*_M + q_M)$ must cover the whole length $1/q_M$. If q_M odd however, as in Figure 7.6, $x_n(\sigma_M)$ will be roughly antipodal to $x_{n+q_M/2}(\sigma)$, so the magnitude of the error need only be half as large.

Let us start with the even case where $\bar{n}^*_M$ is the lowest n for which

$$(n + q_M)\,|\sigma_M - \sigma| > \frac{1}{q_M} \quad \text{for} \quad 2 \mid q_M \tag{7.28}$$

is satisfied. The inequality is strict because there is no integer that can multiply the irrational $|\sigma_M - \sigma|$ and give the rational number $1/q_M$.

In the odd case where error need only be half as large, we may say that $\bar{n}^*_M$ is the smallest n that satisfies

$$(n + q_M)\,|\sigma_M - \sigma| > \frac{1}{2q_M} \quad \text{for} \quad 2 \nmid q_M \tag{7.29}$$

From (6.17) we can say that $|\sigma_M - \sigma| \approx 1/q^2_M a_{M+1}$ for large M and therefore

$$\bar{n}^*_M(\sigma) \overset{M\gg 1}{\approx} \left\{ \begin{matrix} 1 & \text{for } 2 \mid q_M \\ \frac{1}{2} & \text{for } 2 \nmid q_M \end{matrix} \right\} q_M a_{M+1} - q_M. \tag{7.30}$$

It just remains to invoke (6.14) to arrive at a our final approximate formula for this bound,

Ranges of validity for rational approximants

M	1	2	3	4	5	6	7
$\boldsymbol{n^*_M(\sigma, x_0)}$	2	10	84	3958	108 517	1 454 337	423 742 481
$\boldsymbol{\bar{n}^*_M(\sigma)}$	1	32	259	14 304	253 929	25 839 744	818 183 279
Ratio	2	0.313	0.324	0.277	0.427	0.056	0.518

Table 7.1 — Values for $n^*_M(\sigma, x_0)$, our estimate for its upper bound, $\bar{n}^*_M(\sigma)$, and their ratio for different values of M. Value for $n^*(\sigma, x_0)$ was found with exact numerics, as in 7.7, by comparing a series of $\epsilon_n(\sigma_M, x_0)$ with $\epsilon_n(\sigma_{M+1}, x_0)$. The initial condition x_0 is, as elsewhere, as given in (7.33). Note that $\bar{n}^*_M$ provides an underestimate of n^*_M at low M due to the approximation (6.14).

written

$$\bar{n}^*_M(\sigma) \overset{M\gg 1}{\approx} \left\{ \begin{matrix} 1 & \text{for } 2 \mid q_M \\ \frac{1}{2} & \text{for } 2 \nmid q_M \end{matrix} \right\} q_{M+1} - q_M. \tag{7.31}$$

Whenever n reaches n^*_M, $\epsilon_n(\sigma_M, x_0)$, is no longer equal to the true $\epsilon_n(\sigma, x_0)$, and to be certain one is obtaining true sequences based on $\epsilon(\sigma)$ using $\epsilon(\sigma_M)$ one must switch from M to $M+1$ and approximate it with a higher degree of precision. This means insight into the true sequences can be gained by seeing how the approximants change with increasing M.

In what follows, I will use the bound $\bar{n}^*$ as an estimate for the exact value of $n^*_M(\sigma)$. Figure 7.7 demonstrates a fairly good agreement between the estimate, (7.31), and the numerical values, even including the prefactor of $1/2$, as shown in Table 7.1. These values are calculated with our particular σ, x_0 given in the next section. There is a reason for this good fortune owing to our choice of σ, which I will return to in § 7.2.5.

7.2.3 Notes on the numerics

The calculations were done using Mathematica and Python. In both cases, exact fractions were used to calculate $x_n(\sigma_M)$ in equation (7.15) as two integers. This avoids introducing any floating point arithmetic errors, even at very long times. The limitation is on the computing resources. For reasons discussed just following equation (7.20), one cannot define an algorithm for $S^{\Pi}_n(\sigma_M)$ that generates the n^{th} term of the series without knowing the values prior [303] as can be done with $x_n(\sigma_M)$ and therefore $\epsilon_n(\sigma_M)$. Unless some "seed" value for $\epsilon^{\Pi}_n(\sigma_M)$ at some n is known as a starting point, the entire series of $x(\sigma_M)$ and $\epsilon^{\Pi}_n(\sigma_M)$ must be calculated from scratch and held in memory for the calculation. For high values of n ($\sim 10^9$), this means holding two integers of around 15 digits each for each new value of $S^{\Pi}_n(\sigma)$ desired. This can overwhelm the memory of a standard laptop or desktop computer and was the limiting factor in our computational resources.

Although much of the following is devoted to how $S^{\Pi}(\sigma_M)$ behaves with increasing M (similar to how Wang et al. studied the behavior of K with increasing N) truncating this at $M = 9$ gives the approximant

$$\sigma_9 = [1;\, 4,\, 9, \ldots,\, 81] = \frac{170195445997}{136891356233}\,. \tag{7.32}$$

In § 7.2.2 I described how we used two rational sequences to verify that they reproduces the sequence of $\epsilon_n(\sigma)$ values for exact σ, and (7.32) holds good up to $n \approx 10^{11}$, well beyond the

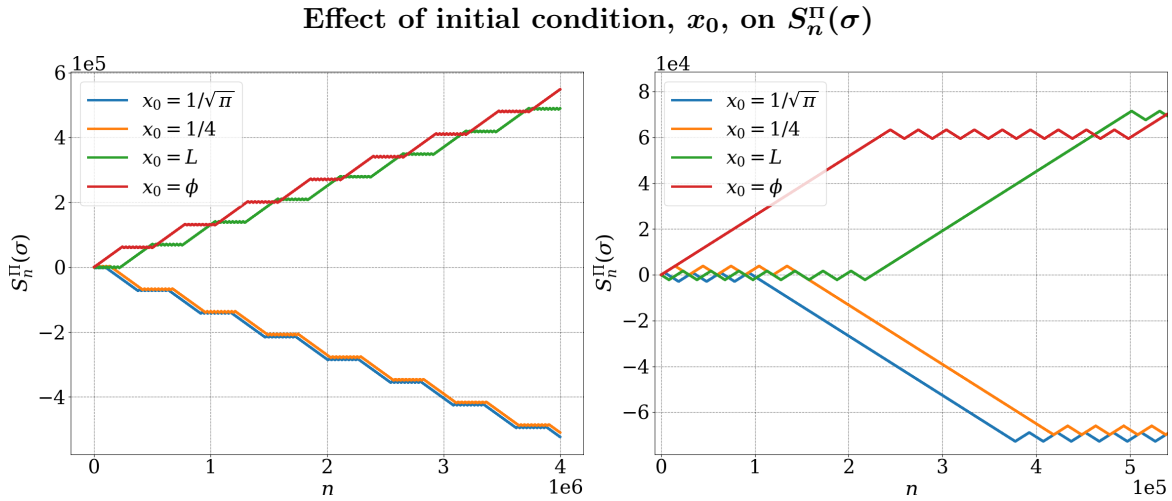

Figure 7.8 — Plots showing the effect of different values for the initial condition, x_0, have on the behavior of $S_n^{\Pi}(\sigma)$. The values are as defined in equation (7.33), 1/4, Liouville's constant, L as defined in equation (7.9), an irrational number very close to a rational number, and ϕ as defined in equation (7.7), an irrational number very far from a rational one. The right subfigure is a blow-up of same series' plotted on the left. See text for exact details.

maximum value used in the following results which only go up to $n \approx 10^9$. This fraction was used whenever we needed to calculate any "true" value or sequence.

The results presented here use a certain initial condition unless otherwise noted. Its value is

$$\begin{aligned} x_0 &= \frac{1}{\sqrt{\pi}} = 0.564189583547756287\ldots \\ &\approx \frac{62840}{111381}\,, \end{aligned} \tag{7.33}$$

with the fractional expression being what was used in the exact arithmetic numerical propagation. Like (7.32), for all practical purposes, this rational approximation is a valid replacement for the exact value.

Ideally, the exact value of the initial condition will be irrelevant to the behavior of an ergodic system. That is one reason that I was so careful to construct $\bar{n}_M^*$ to be independent of x_0. To gather some evidence for this, I checked other values of x_0 to see make sure that the results presented here are not dependent on our choice of x_0. Figure 7.8 indicates the results. It amounts to the expected result of a phase shift in the starting point which in turn can flip the overall sign of $S_n^{\Pi}(\sigma)$, but it does not change the manner of absolute growth and therefore cannot affect implications about the ergodicity. Hereafter, $x_0 = 1/\sqrt{\pi}$ unless otherwise noted and it will be dropped from the notation wherever possible.

The exact values used to make Figure 7.8 are the ratio expression in (7.33), 1/4, Liouville's constant as defined in equation (7.9) with the sum running up to $k = 4$ which made a fraction with numerator and denominator of 29 digits, and $\phi = [0; 1, 1, \ldots, 1]$ with two-hundred ones, which gave a fraction with a numerator and denominator of 42 digits.

I also looked briefly at the effect of other values of the shift, σ. The results are presented in Figure 7.9. It seemed that irrational shifts closer to rationals lead to steeper average slopes of

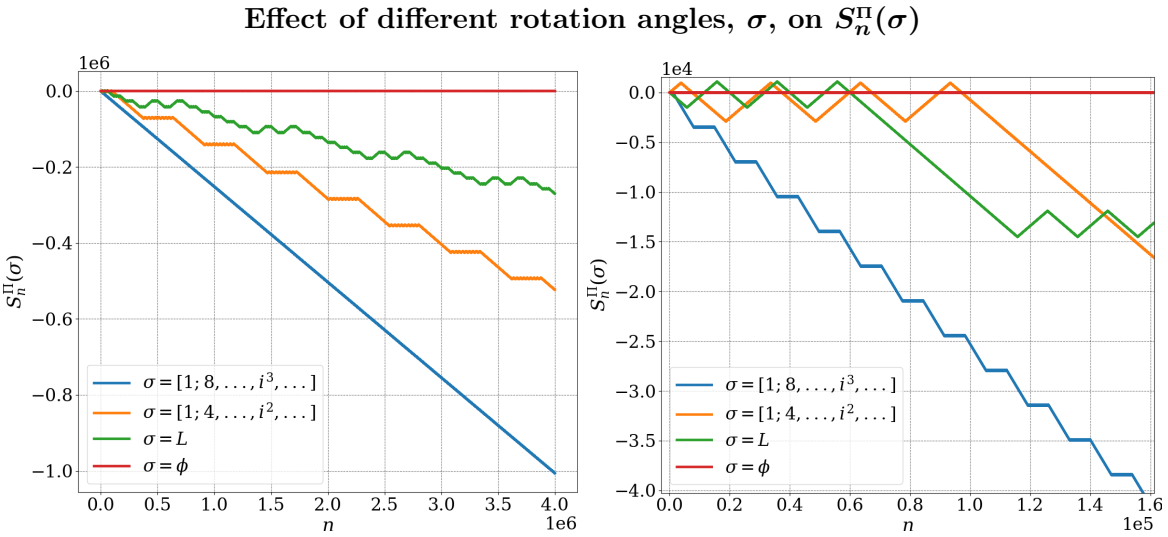

Figure 7.9 — Plots showing the effect of different values for the irrational rotation value, σ, have on the behavior of $S_n^{\Pi}(\sigma)$. The values are as defined in 7.32, a similar generic irrational with $a_i = i^3$, Liouville's constant, L as defined in equation (7.9), and ϕ as defined in equation (7.7). The right subfigure is a blow-up of same series' plotted on the left. See text for exact details.

$S_n^{\Pi}(\sigma)$. Listed in order from "furthest from" to "closest to" the nearest rational, the four values for σ presented in Figure 7.9 are ϕ, the most severely bounded constant type irrational. Then Liouville's constant, L, which is a generic irrational, however, one will always appear in its continued fraction expansion, even at very high indices [300]. This means that comparatively poor approximations exist, in contrast with our studied value of (7.24), where each successive approximation is always better than the one before it. Lastly there is another generic irrational that approaches a rational even faster than (7.24); one where the continued fraction coefficient is the cube of the index instead of the square. Figure 7.9 seems to indicate that the slope of $S_n^{\Pi}(\sigma)$ is steeper with irrational shifts that are better approximated by rationals. This is consistent with our argument for unbounded growth failing on constant type irrationals but would require its own study to say more about.

As before, L is calculated with k running up to four and ϕ from a continued fraction with two-hundred ones. The other two values of σ_M were calculated to a precision of $M = 9$; which, as mentioned before, will give the exact sequence in the range plotted.

7.2.4 Counter growth

With all of this out of the way, I can set down our construction for the unbounded subsequences of $S_n^{\Pi}(\sigma)$.

Consider again the observable f,

$$f(x) = \begin{cases} +1 & \text{for} \quad x \in \left[0, \frac{1}{2}\right[\\ -1 & \text{for} \quad x \in \left[\frac{1}{2}, 1\right[. \end{cases} \tag{7.34}$$

If σ is irrational then the ergodicity of irrational rotations implies that its partial temporal mean, $\bar{f}$ (6.4), converges to its ensemble average, $\langle f \rangle$ (6.2), and equation (6.26) holds true for

(7.34). Huveneers in [259] makes this the increment, $\epsilon_n \equiv f(x)$, of a counter, $S_n(\sigma)$, defined as

$$S_n(\sigma) \equiv n\overline{f}_n(\sigma) = \sum_{n'=0}^{n-1} \epsilon_{n'}(\sigma)\,. \tag{7.35}$$

Ergodicity then means that

$$\langle\epsilon\rangle \equiv \langle\epsilon(\sigma, x_0)\rangle_{x_0} = \langle f\rangle \equiv \int_0^1 f(x)\,\mathrm{d}\mu\,. \tag{7.36}$$

Since the measure, $\mu(x)$, is invariant under the rotational flow, $\mathrm{d}\mu = \mu(\mathrm{d}x) = \mathrm{d}x$ is constant for all x and an integral over it only differs from one over $\mathrm{d}x$ by a factor of the total phase volume which in the case of the Lebesgue measure of the interval $[0, 1[$ is just one. This makes (7.36) read

$$\langle\epsilon\rangle = \int_0^1 f(x)\;\mathrm{d}x = 0\,, \tag{7.37}$$

since $f(x) = -1$ for half the values of x and $+1$ for the other half. The exact same thing can be said for the average over the initial values of $S_n^{\Pi}(\sigma)$, since $e_n^{\Pi}(\sigma, x_0) = -1$ for half the values of x_0 and $+1$ for the other half for any given n, σ. Ergodicity of f, then ensures that the time average of ϵ is also zero. Since S_n is by definition (7.35) the time average of this increment, multiplied by the time step, this means that the counter $S_n(\sigma)$ grows strictly more slowly than linearly with time, n:

$$S_n(\sigma) = o(n)\;.$$

This is a fine example of how to connect the ergodicity of a system to the growth of a counter.

However, in [259], Huveneers goes on to show that S_n is not bounded from above, also important for ergodicity because it would conflict with metric transitivity the way K being bounded from above would. He does this by constructing unbounded subsequences. We use similar methods to construct our unbounded subsequences.

This highlights the two ways we have found to use counters to mount argument about ergodicity. There is the method of relating them to an average and then showing their growth does or does not allow that average to converge, and there is the method of making the counter keep track of states of the system and using its infinite limit to find invariant sets.

In order to asses the growth of $S_n^{\Pi}(\sigma)$ we examined the behavior of the rational $S_n^{\Pi}(\sigma_M)$. The true $S_n^{\Pi}(\sigma)$ sequence will be the same up until n^*. The question then becomes what happens to $S_n^{\Pi}(\sigma_M)$ for increasing M. If we can show that $S_n^{\Pi}(\sigma_M)$ grows, on average, as n and M increase, we have shown $S_n^{\Pi}(\sigma)$ is unbounded and therefore, ergodic.

Periodicity of $x_n(\sigma_M)$ (see equation (6.19)) guarantees the periodicity of $\epsilon_n(\sigma_M)$ with the same period q_M since it is a function of x. This in turn is sufficient to ensure that the parity counter increment, $\epsilon_n^{\mathrm{IR}\Pi}(\sigma_M)$, is either periodic, or anti-periodic, with the same period. Which it is depends on the sign of the step just before it. That is,

$$\epsilon_{n+q_M}^{\Pi}(\sigma_M) = \epsilon_{q_M-1}^{\Pi}(\sigma_M)\epsilon_n^{\Pi}(\sigma_M)\,. \tag{7.38}$$

This leads to two overall patterns when it comes to the sum of these increments. If $\epsilon_{q_M-1}^{\Pi}(\sigma_M) = +1$ then $\epsilon_{n+q_M}^{\Pi}(\sigma_M) = \epsilon_n^{\Pi}(\sigma_M)$ and the whole sequence of increments within

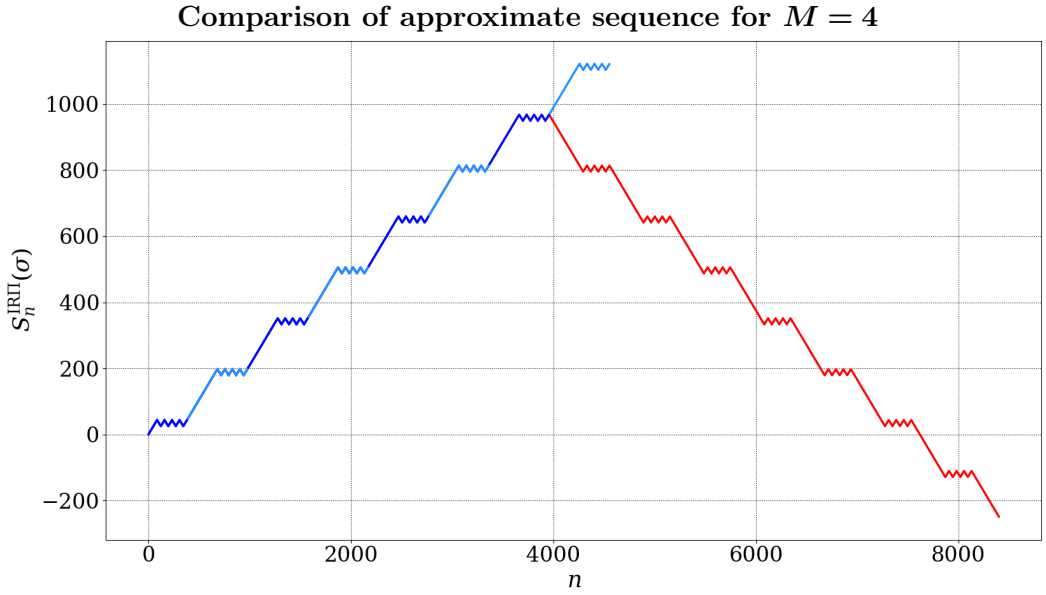

Figure 7.10 — The exact sequence S_n^{Π} (red). The blue curve is the sequence for the convergent σ_M for $M = 4$. Alternating light and dark blues signify periods in the $\epsilon_n(\sigma_4)$ sequence of length $q_4 = 496$. It is plotted from the origin up to one orbit beyond its point of divergence, $n_4^* = 3958$ (see Section 7.2.2). Note that almost immediately after crossing the border of validity, the true sequence stops increasing, while the approximate one continues doing so. This is the S_n^{Π} sequence moving from a "growth" subsequence to a "plateau" subsequence (see text) with longer period and larger amplitude (see Figure 7.11).

any $\Delta n = q_M$ will be the same. Therefore

$$\sum_{n=0}^{q_M} \epsilon_n^{\Pi}(\sigma_M) = \sum_{n=q_M}^{2q_M} \epsilon_n^{\Pi}(\sigma_M)$$

and $S_n^{\Pi}(\sigma_M)$ will grow by steps of $S_{q_M}^{\Pi}(\sigma_M)$, whatever that happens to be, over every time interval of q_M. This means the counter grows linearly, on average, for that rational σ_M. If, on the other hand, $\epsilon_{q_M-1}^{\Pi}(\sigma_M) = -1$ the subsequence of length $\Delta n = q_M$ gets added on with a minus sign and the counter returns to its original value after every other cycle. This forms a plateau of, on average, no growth. Breaking the counter up into these cycles of length q_M we can say that for a rational rotation

$$\begin{aligned} S_n^{\Pi}(\sigma_M) &= \left\{ \begin{array}{lll} N S_{q_M}^{\Pi}(\sigma_M) + S_{\Delta n}^{\Pi}(\sigma_M) & \text{for} & \epsilon_{q_M-1}^{\Pi}(\sigma_M) = +1 \\ S_{\Delta n}^{\Pi}(\sigma_M) & \text{for} & \epsilon_{q_M-1}^{\Pi}(\sigma_M) = -1 \text{ and } N = \text{even} \\ S_{q_M}^{\Pi}(\sigma_M) - S_{\Delta n}^{\Pi}(\sigma_M) & \text{for} & \epsilon_{q_M-1}^{\Pi}(\sigma_M) = -1 \text{ and } N = \text{odd} \end{array} \right\} \\ &\sim \left\{ \begin{array}{ll} N S_{q_M}^{\Pi}(\sigma_M) & \text{for } \epsilon_{q_M-1}^{\Pi}(\sigma_M) = +1 \\ S_{q_M}^{\Pi}(\sigma_M) & \text{for } \epsilon_{q_M-1}^{\Pi}(\sigma_M) = -1 \end{array} \right\} , \end{aligned} \tag{7.39}$$

with N and Δn being defined through Euclidean division of n by q_M:

$$n = N q_M + \Delta n \ , \tag{7.40}$$

which implies

$$0 \leq \Delta n < q_M \ . \tag{7.41}$$

Dividing by q_M is the natural choice since while the behavior is captured by p_M/q_M, q_M is the smallest possible return time [304].

Below, a table for a specific, representative, hypothetical case of $q_M = 3$, and therefore $\epsilon_{n+Nq_M}(\sigma_M) = \epsilon_n(\sigma_M)$ worked out explicitly up to $N = 3$, illustrates this statement:

For this table:
$q_M = 3$
$\epsilon_n \equiv \epsilon_n(\sigma_M)$
$S_n \equiv S_n^{\Pi}(\sigma_M)$
$\epsilon_3 = \epsilon_0$, $\epsilon_4 = \epsilon_1$, $\epsilon_5 = \epsilon_2$

$$
\begin{aligned}
S_0 &= 0 \\
S_1 &= 0 + \epsilon_0 \\
S_2 &= 0 + \epsilon_0 + \epsilon_0\epsilon_1 \\
S_3 &= 0 + \underbrace{\epsilon_0 + \epsilon_0\epsilon_1 + \epsilon_0\epsilon_1\epsilon_2} \\
&= S_3 \\
S_4 &= 0 + \epsilon_0 + \epsilon_0\epsilon_1 + \epsilon_0\epsilon_1\epsilon_2 + \epsilon_0\epsilon_1\epsilon_2\epsilon_0 \\
S_5 &= 0 + \epsilon_0 + \epsilon_0\epsilon_1 + \epsilon_0\epsilon_1\epsilon_2 + \epsilon_0\epsilon_1\epsilon_2\epsilon_0 + \epsilon_0\epsilon_1\epsilon_2\epsilon_0\epsilon_1 \\
S_6 &= 0 + \underbrace{\epsilon_0 + \epsilon_0\epsilon_1 + \epsilon_0\epsilon_1\epsilon_2} + \underbrace{\epsilon_0\epsilon_1\epsilon_2\epsilon_0 + \epsilon_0\epsilon_1\epsilon_2\epsilon_0\epsilon_1 + \epsilon_0\epsilon_1\epsilon_2\epsilon_0\epsilon_1\epsilon_2} \\
&= S_3 + (\epsilon_0\epsilon_1\epsilon_2) \times S_3 \\
S_7 &= \dots \\
&\dots \\
S_9 &= S_3 + (\epsilon_0\epsilon_1\epsilon_2) \times S_3 + (\epsilon_0\epsilon_1\epsilon_2)^2 \times S_3 \\
&\dots
\end{aligned}
$$

Now let us consider what this means for the irrational case where we must change M to $M+1$ every time n reaches $n^*(\sigma_M)$. While $\epsilon^{\Pi}_{q_M-1}(\sigma_M) = +1$, and $n < n^*(\sigma_M)$, $S^{\Pi}_n(\sigma) = S^{\Pi}_n(\sigma_M)$ and will grow linearly. Once n reaches n^* we must increase M to $M' = M+1$ and, in the worst case, $\epsilon^{\Pi}_{q_{M'}-1}(\sigma_{M'}) = -1$ and we enter a plateau subsequence. This will, on average, preserve the absolute value the counter gained during its growth period. The other case, would be that it entered another growing subsequence. In either case, it will not head back towards zero unless it is to cross it over the course of a cycle. Since n^* can always be increased by raising M, this pattern can be expected to continue for as long as there exist higher values of M for which $\epsilon^{\Pi}_{q_M-1}(\sigma_M) = +1$. If there were no such values, then one could say that, after some n, irrational rotations would *always* put an odd number of points in the interval $[1/2, 1[$ since (7.39) holds for all n. This cannot be since irrational rotations are known to be ergodic. Therefore, $S^{\Pi}_n(\sigma)$ should be unbounded. Overall, we would conjecture that this counter, $|S^{\Pi}_n(\sigma)|$, grows like n on average (see 7.13), something my collaborators and I would like to prove in future work. Figure 7.13 indeed shows the overall growth of the counter in question.

This is the key result because it is true for any parity counter. The goal is to construct other parity counters, such as (7.14), and examine them and this method could be applied wherever we can construct a parity counter and find reliable cycles in it.

Note that this whole line of reasoning requires that the rational approximant hold for at least two orbits of σ_M when $\epsilon^{\Pi}_{q_M-1}(\sigma_M) = -1$. If n^* were to be reached before $S^{\Pi}(\sigma)$ completed

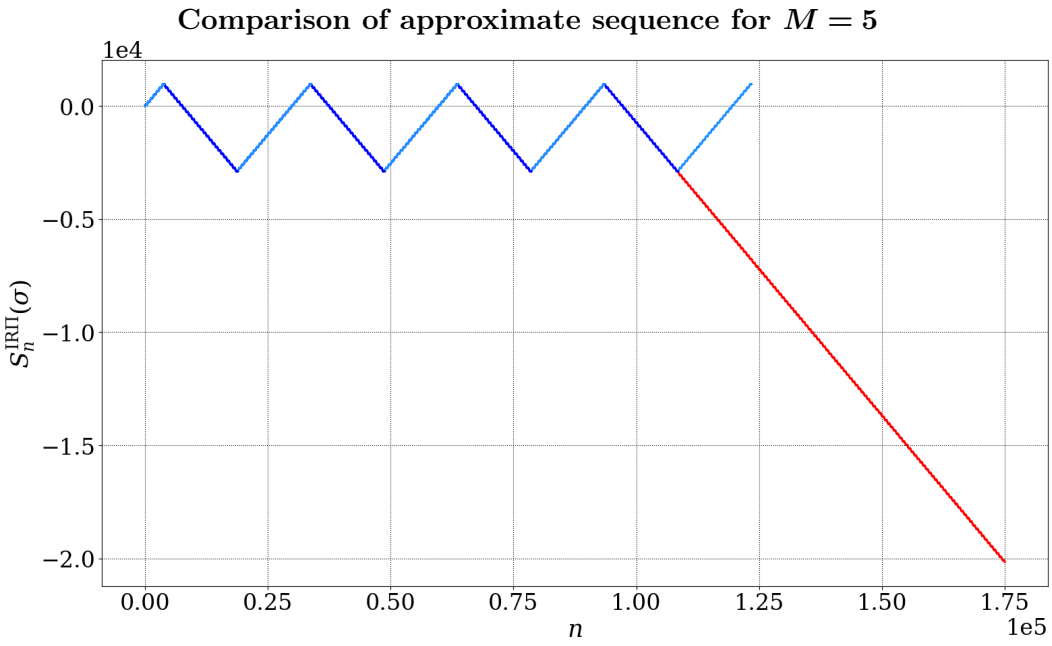

Figure 7.11 — The same as Figure 7.10, but for approximant sequence $M = 5$. Here we can clearly see the plateau cycle with amplitude ≈ 5000 associated with $M = 5$ followed by a growth cycle beginning after $n_5^* = 108\,517$. Note that the ramps in the plateau subsequence, as well as the growth phase of the true sequence, are not straight lines but "stepped" as they are in Figure 7.10. This is the clue to the graph's fractal nature.

the cycles predicted in equation (7.39), we cannot say that the counter returns to the value at which it began these cycles and the absolute value gained in preceding growth phases may be wiped out. Also, since growth phases may differ in sign it is important that arbitrary large ones may always be found with increasing M to ensure the growth of the absolute value. This is why we insist that the irrational be of generic type, since for a constant type irrational there is some maximum value for the continued fraction coefficients and therefore an upper limit to the number of cycles that can be expected in a phase. To illustrate this, I rewrite the estimate for n^* as

$$\bar{n}_M^*(\sigma) \overset{M\gg 1}{\approx} \left\{ \begin{matrix} 1 & \text{for } 2 \mid q_M \\ \frac{1}{2} & \text{for } 2 \nmid q_M \end{matrix} \right\} \frac{1}{|\sigma_M - \sigma| q_M} \,. \tag{7.42}$$

From equation (6.15), $|\sigma_M - \sigma|$ has an upper bound of $1/q_M^2 a_{M+1}$ therefore $\bar{n}_M^*(\sigma) \gtrapprox q_M a_{M+1}$, neglecting the factor of two for the different cases, and one cannot expect more than a_{M+1} cycles of length q_M to hold for a given M. If σ is of constant type, this has a maximum value and in that case, since the growth phases are of uncertain sign and may cancel out, we cannot guarantee unbounded growth.

7.2.5 Specifics of our choice of irrational

We constructed an irrational whose by saying that the continued fraction coefficients are their index raised to a power. This guarantees it is generic but it has properties that might make it less representative.

Firstly with regard to growth, we can say something further than what was set down in the last section; although it does not rise to a model for the overall growth. These growing

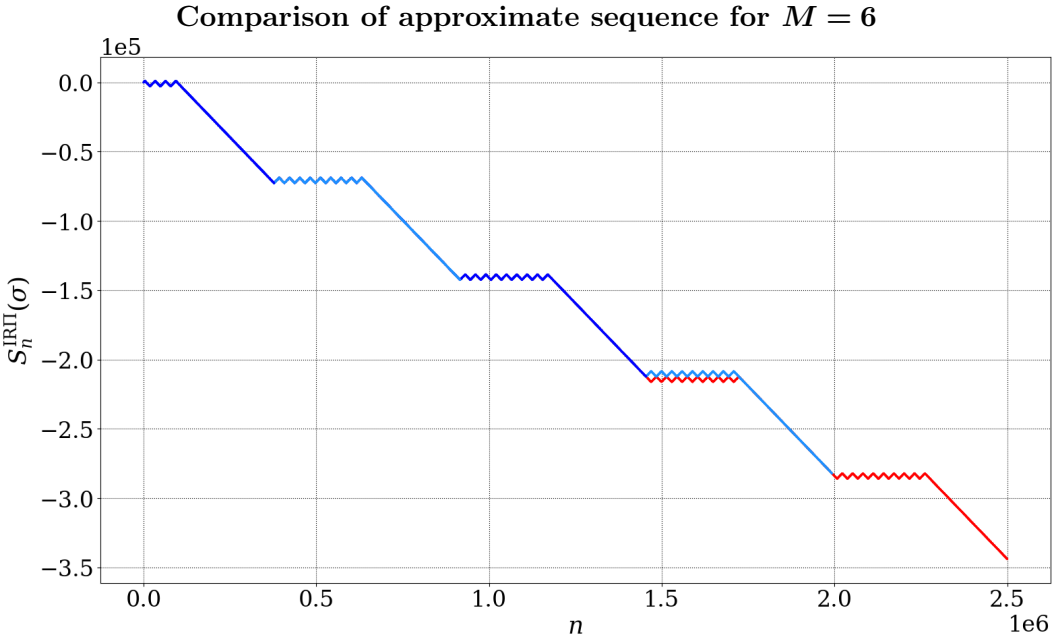

Figure 7.12 — The same as Figure 7.10, but for $M = 6$. Here we can see that *both* $M = 6$ *and* $M = 7$ correspond to growth cycles even though they get out of sync at $n_6^* = 1\,454\,337$. This is proof of the statement made in the text that growth and plateau cycles need not alternate. Note the plateau phases of amplitude ≈ 5000 as could be seen in Figure 7.11.

subsequences in $S_n^{\Pi}(\sigma_M)$ are in fact guaranteed on alternating values of M when a_M is M raised to an integer power.

If q_M is a multiple of four, then $\epsilon^{\Pi}_{q_M-1}(\sigma_M) = +1$ and $S_n^{\Pi}(\sigma_M)$ will grow as in the first case in (7.39). This is a result of the fact that multiples of four are doubly even, so when the q_M points are laid equidistant on the interval $[0, 1[$ an even number of them will fall in the interval $[1/2, 1[$. This will give an even number of instances of $\epsilon(\sigma_M) = -1$ which will give $\epsilon^{\Pi}_{q_M-1}(\sigma_M) = +1$.

Next, let us show that for $a_M = M^i$ for an integer $i \geq 2$, $q_M \equiv 0, 1 \mod 4$, with $q_M \equiv 0 \mod 4$ when M is even and $q_M \equiv 1 \mod 4$ when M is odd in our numbering convention, (7.21). First consider the even case: $M = 2k$ where $k \in \mathbb{N}$. From (6.13) we write

$$q_{2k} = (2k)^i q_{2k-1} + q_{2k-2}\,.$$

Notice that $(2k)^i$ contains a factor of $2^2 = 4\,\forall k > 0$, so the first term on the right hand side will be $0 \mod 4$ wherever q_M is well defined, i.e. $k \geq 2$. The last term then implies that all $q_M \equiv q_{M'} \mod 4$ for M, M' both even. Since $q_2 = a_2 = 2^i$ will always contain a factor of 4 for $i \geq 2$, $q_{2k} \equiv 0 \mod 4\ \forall k$. This means that whenever M is even $S_n^{\Pi}(\sigma_M)$ will be in a linear growth subsequence as described in the first line of (7.39).

For the odd M case (6.13) gives

$$q_{2k+1} = (2k+1)^i q_{2k} + q_{2k-1}\,.$$

Since we have shown that $q_{2k} \equiv 0 \mod 4$, the above implies that $q_M \equiv q_{M'} \mod 4$ for M, M' both odd as well. Since $q_1 = 1\ \forall i$, $q_M \equiv 1 \mod 4$ for odd M. This does *not* imply however

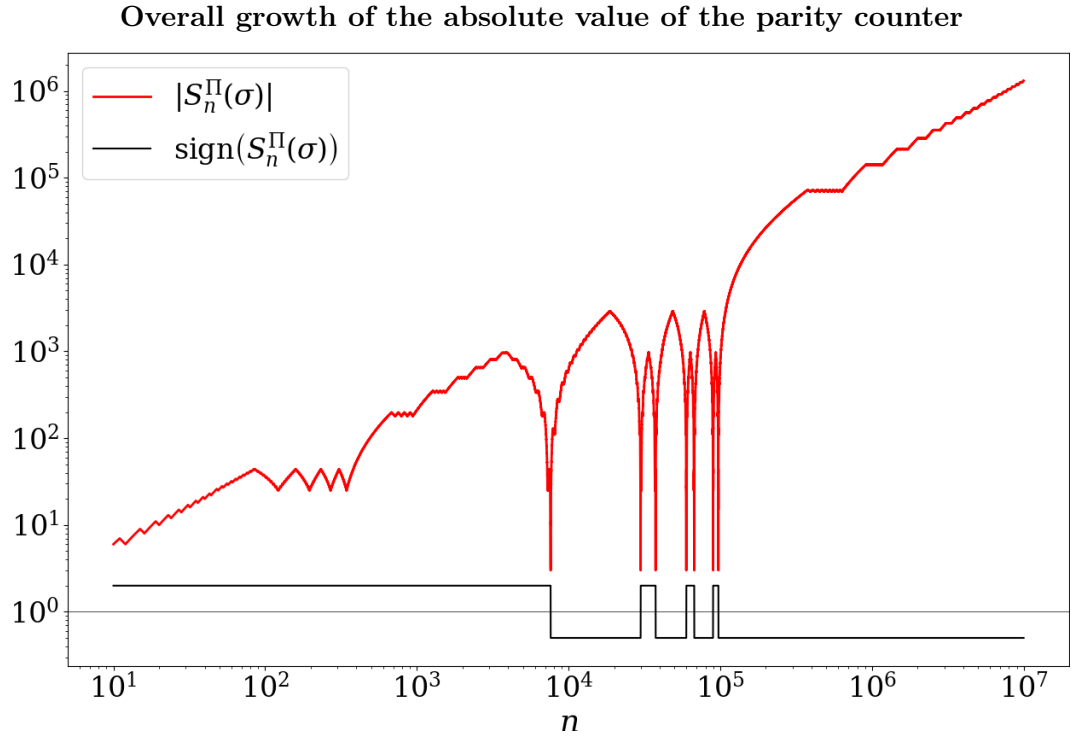

Figure 7.13 — The absolute value of the exact S_n^{Π} sequence up to $n = 10^7$ is show in red. The deep valleys do, in fact, fall to zero, but I chose a minimum value of 3 for the sake of the plot. The sign of S_n^{Π} is shown in black, with 2 on the log scale (above the horizontal rule) indicating $+1$ and $1/2$ (below it) indicating -1. Notice that while there are zero-crossings, the absolute value is not localized near zero.

that $\epsilon^{\Pi}_{q_M-1}(\sigma_M) = -1$ for odd M and that S_n^{Π} switches from a linear growth interval to a "plateau" interval when M switches from even to odd. It guarantees that there will be an unpaired point every other time $M \to M+1$, and so a chance to switch, but that extra point may fall in the interval $[0, 1/2[$ and so change nothing. Whether this happens for any particular M will depend heavily on the initial condition, x_0, but the growth of the counter only depends on how frequently it happens, which has nothing to do with x_0.

Further, we can say avoid a pitfall for $2 \mid q_M$. If $\sigma < \sigma_M$ for $2 \mid q_M$ and $x_0 = 0$, then $x_{q_M/2}(\sigma_M, 0) = 1/2$ implying $\epsilon_{q_M/2}(\sigma_M) = -1$ but $\sigma < \sigma_M$ means that $\epsilon_{q_M/2}(\sigma) = 1$ and we have arrived an n^*_M before completing even one period. But when the continued fraction coefficients are their index raised to a power, that cannot happen since $2 \mid q_M$ if and only if $\sigma_M < \sigma$. Every convergent, that is with $2 \mid M$, is smaller than σ [294, Trm. 154]. As we have shown above, these even convergents are exactly those where $2 \mid q_M$ and so we avoid this problem. This is not a general property of generic irrationals and will have to be dealt with when generalizing these methods. But when define other counters for other systems we will have to find the periodicity in *those* systems and that will involve constructing another way to estimate the range of validity.

7.2.6 Fractal nature

As noted in the last section, for any given M, the behavior of $S_n^{\Pi}(\sigma_M)$ will contain the behavior of all of the quasi-periods corresponding to proper periods in $\epsilon_n^{\Pi}(\sigma_{M'})$ for $M' < M$

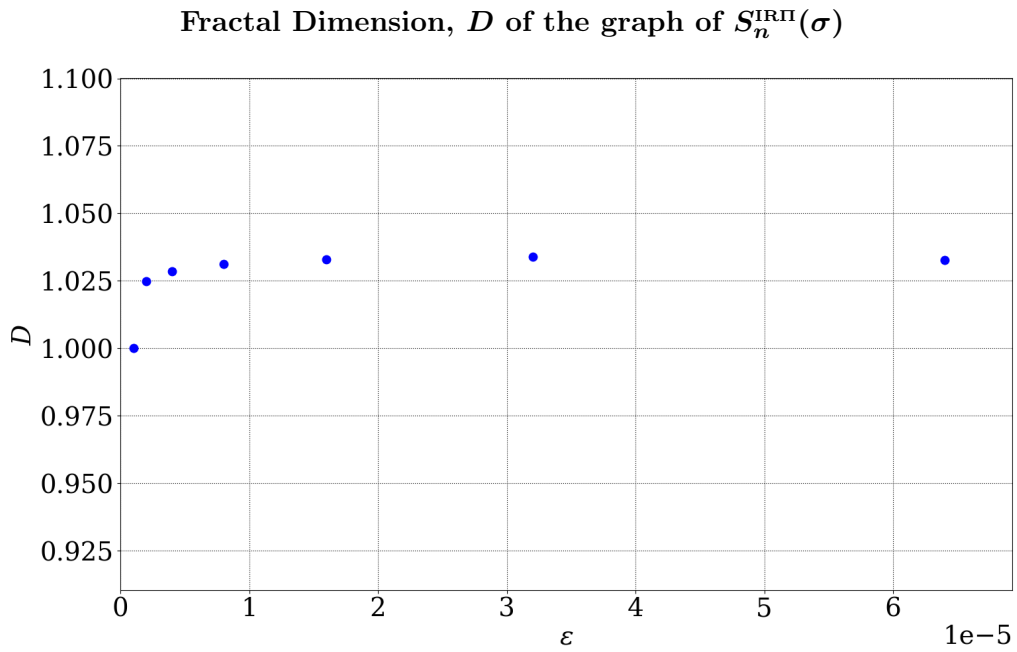

Figure 7.14 — Estimated fractal dimension D of the S_n^{Π} curve as a function of the course-graining parameter ε (see text and equation (7.45)). The smallest ε corresponds to the exact curve. As the ε increases, and more microscopic details of the curve become blurred, the dimension quickly saturates at $D \approx 1.03$.

(see Figures 7.10 through 7.12). This gives a self-similar nature to $S_n^{\Pi}(\sigma)$, which must contain *all* of these quasi-periods, and implies that its plot might have a fractal dimension.

Visually, simply looking at Figures 7.10 through 7.12 together indicates a fractal nature for the S_n^{Π} curve. Indeed, both the growth intervals,

$$0 \leq n < n_M^*, \quad M : \epsilon_{q_M-1}^{\Pi}(\sigma) = +1, \tag{7.43}$$

and the plateau intervals,

$$0 \leq n < n_M^*, \quad M : \epsilon_{q_M-1}^{\Pi}(\sigma) = -1, \tag{7.44}$$

include all the previous intervals as building blocks. Calculating the Minkowski (i.e. box counting) dimension [305, §3.1] bears this out as can be seen in Figure 7.14 where the fractal dimension of the S_n^{Π} graph appears to saturate close to $D \approx 1.03$.

Our method of computing the fractal dimension is one given in [306] for a one-dimensional curve in two-dimensional space: first imagine that the curve described by S_n^{Π} is drawn over a 2D rectangle of size $n_{\max} \times S_{\max}$. This space is then divided onto squares of dimension $\varepsilon(n_{\max}+1) \times \varepsilon(n_{\max}+1)$. The case of $\varepsilon = (n_{\max}+1)^{-1}$ would capture curve exactly since the curve is drawn in $n_{\max}$ discrete steps. Then, the number of the boxes that contain the curve, $\mathcal{N}_{\text{filled}}$, is computed. For a given ε, the estimate of the fractal dimension reads

$$D = \frac{\ln(\mathcal{N}_{\text{filled}})}{\ln(1/\varepsilon)} \ . \tag{7.45}$$

Since the curve is in fact a discrete set of points, one for each n, $D_0 \equiv \lim_{\epsilon \to 0} D = 1$ as can be seen in Figure 7.14. But it quickly saturates to $D \approx 1.033$ around $\varepsilon = 5 \times 10^{-6}$ and it appears to the grid as a continuous curve with non-integer dimension.

7.3 Conclusion

In this chapter I detailed the WCP counter, K, along with its underlying 1-D system and billiard. I described the conflict the authors see between the localization around the initial value, K_0, of their counter and the fact that all values of K are accessible from any K_0 given enough time: that appears to be a non-ergodic system with an undivided phase space, in conflict with what KAM theory finds for all known non-ergodic systems [70]. I took their advice about numerical results and examined exactly how strongly these results conflict with the billiard being ergodic and find that the two are not irreconcilable. An ergodic billiard could produce their results and if their counter is unbounded as they say, it would suggest that the billiard is metrically transitive (at least in its velocity angle) and therefore ergodic.

This serves a counter constructed by M. Olshanii which counts the parity of $f = -1$ events in a rotation. The hope is that the methods for examining its growth could be used for other parity counters such has (7.14), and give insight into the WCP system as well as other billiards. The two lines of attack it opens are direct arguments about the convergence of averages as at the beginning of § 7.2.4 or arguments involving unboundedness and metric transitivity. I give a detailed description of the methods we used to examine it via rational approximants. I defined and gave ways of estimating up to what time a given approximant will be valid. When the irrational is generic and therefore can be well approximated, I explained our proof demonstrating that the counter contains unbounded subsequences and therefore ergodic. Because the overall sequence is made up of nested subsequences, it should have something of a fractal nature and I present our calculation for the fractal dimension of its graph, $D \approx 1.03$. The fact that this dimension is not an integer lends weight to the idea that the self-similarity we see in Figures 7.10 through 7.12 is not an artifact. While shapes with non-integer fractal dimension are not necessarily self-similar, self-similarity is a way for a shape to acquire a non-integer fractal dimension [307]. Since we have a strong argument why these curves should be self-similar (as set out in (7.39)), it stands to reason that the non-integer fractal dimension comes from that rather than that the self-similarity is illusory and it acquires a non-integer fractal dimension from some other source. I also described the numerical work I did to make sure that the values we used for the shift σ and the initial condition, x_0, were not special and that varying the values for these quantities did not lead to any surprises.

The next step is to define other parity counters over other systems and apply methods such as the cycles described in (7.39) to them. The ergodicity of irrational rotations is already known. The natural choice is a billiard of some kind. We think that with slight modifications, we will be able to build bridges from what we have learned here to certain billiards. For example, by modifying f as written in (7.34) to have the form

$$f' = \begin{cases} +1 & \text{for} \quad x < \frac{1}{2} + \delta \\ -1 & \text{otherwise,} \end{cases} \tag{7.46}$$

where δ is a positive or negative shift that depends on the trajectory, one moves closer to—without actually arriving at—the case of a $45° - 45° - 90°$ triangular billiard. Another idea is to consider an irrational rotation by angle σ along a circle of unit circumference. Then instead of being divided in half, the circle is divided into two parts, one of length σ and the other of $1 - \sigma$. After every rotation, the counter increments by a step, $\epsilon = \pm 1$, depending on the parity of the prior number of landings to the $1 - \sigma$ zone. We suspect that if we can show that this contains an unbounded subsequence, we can argue that the WCP counter is also

unbounded. We made the connection between the WCP counter in a $45° - 45° - 90°$ triangular billiard and a parity counter on irrational rotation in a very recent paper [308]. In it we show that the counter is unbounded in this case even though such a billiard is not ergodic. This does not imply the converse because the billiard is rational and multiple values of the counter can correspond to the same angle. This means the counter running off to infinity does not mean that the set of angles visited grows at all. However, the important result of a bridge between parity counters on irrational rotations and billiard is detailed, as well as a method for constructing unbounded subsequences beginning from an explicitly defined set of initial conditions.

The next step is to make connections to both irrational right-triangular billiards [309–311], and elongated non-right-triangular billiards as in [312, 313], by constructing other counters.

Conclusion

This thesis presents preparatory work for two methods of examining one-dimensional systems, one experimental and one theoretical. At the conclusion of this work both are in a similar state: almost ready to be put to use to obtain results. The experiment needs a little more work to be able to load the atoms onto the chip as described in Chapter 3. The current profiles for the various stages of loading must be implemented, imaging systems tested and installed, etc. But the chip is in place and under vacuum, and the atoms can be efficiently cooled and transported to the chip.

A key feature of the growth of parity counters (7.39) has been introduced which will hold for all parity counters. It remains to define parity counters for systems other than irrational rotations (whose ergodicity is already known) and look for cycles in them.

I will summarize the results in their respective contexts separately and then make what connections I can between the two.

State of the experiment

The experimental work picks up where D. Ben Ali's thesis work left off [1], bringing the experiment very close gathering data. The chip has been completed and installed and put under a vacuum of $< 10^{-10}$ mbar. The transport has been completed, optimized, and tested up to the science chamber and it performs as expected. Clouds of $\sim 100\,\mu$K can be generated and transported to the science chamber in around ten to twelve seconds at our current oven temperature. Around ten of those seconds are spent loading the MOT and if we increased the temperature of our ovens and increased the atomic flux in our beam, we could cut this down below the transport time of 1.2 s. We expect that this time could be reduced by several seconds, if necessary, by working at a higher oven temperature. Work is now proceeding on implementing the loading and cooling procedures for the chip.

Electrical tests on the chip performed under vacuum indicate that it will be able to generate the fields we need to perform our experiment. I subjected the chip to currents higher than those needed for the experiment for longer duration than the will be used for and it performed as expected. Several months later, T. Badr performed similar tests and observed no change in the chips performance. This suggests that we are not in danger of damaging the chip, or its connections during normal use. The coil used for evaporative cooling, on the other hand, is significantly hampered by being wound around copper. This requires tests with the atoms to see if it can perform adequately.

Parts for the lenses have been obtained and the low-magnification objective has been assembled and is undergoing testing with a clear-optical path target and laser illumination. Some machining remains to be done for the higher-magnification objectives: punching out the apertures for the different imaging conditions and threading the SM2E60 tube. In retrospect, tapping the threads might have been avoided by replacing the SM2E60 6" tube with two SM2L30C, 3" tubes.

The magnetic transport has been optimized to transport the cloud from the MOT chamber to the chip in about 1.2 s while keeping about 80% of the atoms as opposed to three or four times that on comparable experiments [2, 239, 240]. This is done using an error function profile, tuned to be close to a constant velocity profile (5.23). The error function profile, however is smooth in all its derivatives. Classical simulations show us that this profile provides a good compromise between a constant-velocity profile and constant-acceleration profile. The constant-velocity profile necessarily has infinite acceleration spikes at the beginning and end, shifting atoms far from the trap center and leading to finite-depth losses. Capping the acceleration and making a constant-acceleration profile keeps atoms near the center where Majorana spin-flip losses are more prevalent. The error-function profile provides a good compromise.

Perspectives and future work

The chip itself is a combination of tested technologies. Other atom chip experiments have included multiple trapping wires [145] or rf-wires which could dressing the trap [146,219]. Even coplanar waveguides on chips are not a novel [314]. The new territory we intend to explore is the use of a novel Feshbach resonance [43] to reach new parts of the one-dimensional phase diagram. This is the first next step in the experiment once the atoms can be loaded into the trap on the chip: experimentally find the microwave induced Feshbach resonance. The next step will be to use it to explore the physics of the one-dimensional gas. At some point this will entail installing the equipment necessary to perform Bragg spectroscopy and the high-resolution light-sheet imaging system. It will be after that that we intend to explore things like using the trapping wires to excite transverse modes, rf-dressing, or considering the predictions of Yao et al. [191] about the behavior of Tan's contact for a harmonically trapped, strongly-interacting gas.

Further on still, we hope to install optics that will allow us to make a finite-length cylindrical trap. This will be exactly the system described by Lieb and Liniger [48, 49]: particles restricted to one dimension in which they are confined by hard walls. Then, there will be no need to make local density approximations, examine the gas slice by slice, or account for the gas being in different regimes in different places. Since one of the attractive features about performing experiments on one-dimensional systems is the exact solutions to make comparisons with, it would be best to have the precise system which is exactly solved to make comparisons to.

We say that we have optimized our magnetic transport but that requires some qualification. We do not have any theoretical description which allows us to say that we have found a global optimum. The methods for optimizing transport in harmonic and higher-order polynomial traps [242, 244, 245] cannot be used in our case because there is no well-defined derivative at the center of our trap. If further work is to be done optimizing transport in quadrupole trap chains, one can either attempt to build a theoretical description of the transport upon which theoretical tools may then be deployed, or continue in our vein and try to imagine profiles that keep atoms away from the two regions where they would be lost: the edge and center of the trap.

Theoretical Results

The second part is a different way to look at one-dimensional systems. It was inspired by a study using a counter to look at the ergodicity of irrational, right-triangular billiards, which can encode systems of three particles moving in one dimension. They present numerical evidence that this counter implies that the billiard is not ergodic by showing that it localizes. I argue that the billiard might still be ergodic, but a localized counter defined over an ergodic billiard is still an oddity; one expects to see things defined on ergodic systems take on all their available values, like ergodic observables. But counters count events or outcomes, they are not observables themselves.

This localizing counter can be re-cast as a parity counter, one that incitements by ± 1 depending on the parity of the number of some event or occurrences of an observable result. My coauthors an I hope to define other counters to study other systems but we examined the behavior for a counter defined on an irrational rotation to gain insight into how these counters will behave.

The key result is what happens when there is a period in the observable underlying the counter. Once that underlying cycle is found, the absolute value of the parity counter will undergo cycles of average growth or average stasis, depending on the sign of the parity counter increment just before the cycle closes. This is stated in (7.39). If these rational numbers are approximating irrational ones, these cycles will hold until that approximant fails and the next one takes over. This will build up a nest of growth and stasis cycles and estimates of the overall growth can be made, so long as the approximants hold for long enough (two periods). The irrational rotation we chose to study had several convenient features such as alternating cycles of growth and stasis and approximants not failing on $n = q_M/2$ which make it less representative of all irrational rotations, but this is not disastrous since the goal is to define parity counters for other systems.

Perspectives

When generalizing to other systems, such as irrational billiards, we still expect to find trajectories governed by irrational angles. Almost all irrational angles will be generic [294, Thrm. 196], having an unbounded sequence for their continued fraction coefficients. This means that at some point they can be approximated arbitrarily well by that finite continued fraction and cycles in the observable of length q_M will persist for a long time. Look at the upper bound in (6.15) which is inversely proportional to the *next* continued expansion coefficient. Even if the coefficients do not grow monotonically as they do in our test case, one can pick an M at which to truncate the fraction which makes a_{M+1} large compared to q_M. This means that the rational behavior can be expected to hold for a long time (see (7.30) for the case of irrational rotations) and the growth for these periods will be predictable. In the case of generic irrationals, these predictable patches can always be found as one extends the approximant further and further. This opens up two ways to argue for ergodicity: if there are enough predictable patches, it may be possible to find the growth rate and relate it to the mean of a particular observable, thereby showing that the mean converges. Or by showing that the counter is unbounded, showing what that implies about the invariant sets, and then mounting an argument from metric transitivity. Neither of these follows immediately from (7.39) and will depend on the specifics of the system and the counter, but the fact that (7.39) applies to *any* parity counter and that large patches of such cycles can be found for *any* generic irrational means that it should be broadly useful, if appropriate counters can be concocted.

For an example of a parity counter defined on irrational rotations that gives the value of the WCP counter in an integrable billiard, see [308]. This somewhat ameliorates the fact that the irrational we chose to study (7.24) had several convenient features as discussed in § 7.2.5.

Sadly, this insight will be less useful for bounded-type irrationals where there is some limit on the coefficients (7.23). In this case, at some point, q_M will certainly grow large as compared to a_M and the approximant will always fail before the predictable cycle completes. This is certain in the case of irrational rotations and in the case of billiards, where the set of reachable angles grows like q_M. I do not have a neat formula for when the orbit closes which Wang et al. called $\tau^*(q, q')$ and found numerically. They found, as expected, that $\tau^*(q_M, qM + 1)$ grew much slower than q_M for the strongly irrational case where $\alpha/\pi = [0; 1, 1, 1, \ldots]$. In these situations, our method could not be used because cycles would not be stable for long enough.

Bringing these results back to systems like the ones studied in Part I is some ways of. The next step for these counters is to apply them to billiards; constructing things like (7.14) for the WCP billiard and other counters for irrational right triangles, isosceles right triangles, and elongated non-right triangles [312, 313], in order to prove the ergodicity of their flows. But all of these would encode systems of only two particles[2] (and except for the isosceles right triangle) of different masses. The ergodicity of these mass mixtures would not have direct implications for one-dimensional, many-body quantum systems of identical particles. If the solution to the isosceles right triangle generalizes to billiards in higher dimensions, then it might provide a way to show systems like the ones made on the atoms chip are ergodic, at least in the cylinder trap with uniform density. This might be combined with the technique of embedding the billiard [69] to extend it to the case of harmonic longitudinal confinement, but that would only be applicable far in the strongly interacting regime, where the particles are hard-cored.

2. Systems of three particles are also possible using a more complicated coordinate transform. See [95, 250, 285] and references therein.

Appendices

Appendix A

Tables of notations and symbols

General notations

Constants:

Symbol	Description	Value
($\hbar$) h	(reduced) Planck constant	($1.054 \times 10^{-34}\,\mathrm{J \cdot s^{-1}}$) $6.626 \times 10^{-34}\,\mathrm{J \cdot s^{-1}}$
c	Speed of light	$2.998 \times 10^{8}\,\mathrm{m \cdot s^{-1}}$
k_B	Boltzmann constant	$1.381 \times 10^{-23}\,\mathrm{J \cdot K^{-1}}$
μ_B	Bohr magneton	$9.274 \times 10^{-24}\,\mathrm{J \cdot T^{-1}}$ ($h \times 1.399\,\mathrm{MHz \cdot G^{-1}}$)
m	Mass of a sodium atom	$3.818 \times 10^{-26}\,\mathrm{kg}$
a_0	Bohr radius	$5.29 \times 10^{-11}\,\mathrm{m}$

Textual abbreviations:

Symbol	Description
Cf., i.e., e.g.	*Confer*—"Compare,", *id est*—"that is", *exempli gratia*—"for example"
Ch(s)., App., §, p(p).	Chapter(s), appendix, section or subsection, page(s)
f(f).	"and the following (pages, sections, chapters, etc)." Two indicates multiple units of text are relevant
Tbl., Fig.	Table, figure
Eqn., Thrm., Prop., Lm., Pblm. Ex.	Equation, theorem, proposition, lemma, problem, example or exercise

Mathematical notations:

Symbol	Description
$\mathbb{Z}, \mathbb{Q}, \mathbb{R}, \mathbb{N}$	Set of integers, rationals, reals, natural numbers
$\mathbb{S}^+$	Set $\mathbb{S}$ includes only positive numbers
$\check{\boldsymbol{v}}$	A unit vector pointing along the vector $\boldsymbol{v}$
$p \perp q$	p and q are coprime
$p \mid (\nmid) q$	p does (not) divide q
$x \to a^{+/-}$	x goes to a from above/below
$f = O(g)$	$\lvert f\rvert$ Is asymptotically bounded from above by g, i.e. $\limsup_{n\to\infty} \frac{\lvert f(n)\rvert}{g(n)} < \infty$
$f = o(g)$	f Is asymptotically dominated by g, i.e. $\lim_{n\to\infty} \frac{f(n)}{g(n)} = 0$
$f = \Omega(g)$	$\lvert f\rvert$ Is not asymptotically dominated by g, i.e. $\limsup_{n\to\infty} \left\lvert \frac{f(n)}{g(n)} \right\rvert > 0$
$A\ (\subset) \subseteq X$,	A is a (proper) subset of X
$A \setminus X$	Set difference of A and X, i.e. the set those elements in A, but not X

Part I: Sodium experiment

Bose gas properties:

Symbol	Description	Page
$n(\boldsymbol{r})$	Number density at $\boldsymbol{r}$. The one-dimensional analog is $n_{1\mathrm{D}}$	
μ	Chemical potential	
λ_T	Thermal de Broglie wavelength	21
$a_{1\mathrm{D}}$	One-dimensional analog of the 3-D scattering length, a	27
ξ	Healing length	23
$g_m(z)$	Polylogarithm (aka Bose function) of z of order m	20
ℓ_ϕ	Phase correlation length	35
$g_{n\mathrm{D}}$	Contact interaction strength in n-dimensions	22,27
γ	Lieb–Liniger interaction parameter	29
τ	Reduced temperature	34
g^n	Correlation function of order n	32,32
Υ_{Maj}	Spin adiabaticity condition	54

Magnetic trap & transport properties:

Symbol	Description	Page
κ	Quadrupole trap parameter	54
R	Radius of average density in a quadrupole trap	103
r_t	Distance traveled by a particle at average velocity in the trap, v_0	112
C_n	The n^{th} coil in the magnetic transport	109
α	Horizontal anisotropy with $\alpha = 1$ being symmetric	106
δ	Shape parameter for erf & tanh profiles. $\delta \to 0$ being more linear and $\delta \to \infty$ more step-function-like	111
γ	Coefficient to ensure $\dddot{y}_0(\Delta t/2) \to 0$	111
γ_c	Collision rate	114
r_N^2	Atom number ratio, after two-way transport	117
r_T^2	Temperature ratio, after two-way transport	117
r_{Maj}	Parameter to estimate Majorana losses	119

Objective lens design:

Symbol	Description	Page
NA_{obj}, NA_{img}	Numerical aperture (NA) in object or image space	84
M	Magnification, the ratio of image size to object size	85
y/y, θ/θ'	Length or angle in object/image space, unprimed or roman quantities are in object space while primed or italic ones are in image space	84
ν_c	Cutoff frequency for the modulation transfer function (MTF)	85
l_{px}	Pixel size or *pixel pitch*	86
DoF_g	Geometric depth of field (DOF)	86
α_A	Angular Airy diameter	87
r_A	Airy radius	87
dy_{min}	Size of smallest resolvable object	86, 87
DoF_d	Diffractive DOF	87
$M_{\max}$	Maximum useful magnification	88

Part II: Counters and Ergodicity

Symbol	Description	Page
ℓ	Distance between points in an orbit	139
K	Wang-Casati-Prosen (WCP) counter	144
α	Angle used in the definition of the WCP billiard	143
θ	Velocity angle in the WCP billiard	143
φ	Factor to place θ in correct quadrant $\varphi \in \{\theta_0, -\theta_0, \pi - \theta_0, \theta_0 - \pi\}$	144
ϕ	Smaller golden mean, equal $(\sqrt{5}-1)/2$	146
L	Liouville's constant	146
τ^*	Time at which an approximate WCP sequence fails	148
n_K^*	Number of velocity angles visited by τ	148
θ_{y_1}, θ_{y_2}	Angle subtended by the other leg when the WCP virtual particle meets a leg at y_1 (or y_2)	144
$P(K,t)$	Probability that the WCP counter takes the value K at time t	150
$\sigma_M = p_M/q_M$	Rational approximant of an irrational angle, σ	137
(x_0), x	(Initial) Position in a rotation's phase space	138
$\mu(x)$	Measure defined on x	
$\epsilon_n^{\mathrm{IR}} \equiv f_n(x)$	Increment of the irrational rotation counter, $S_n^{\mathrm{IR}}(\sigma, x_0)$	140
$\epsilon_n^{\mathrm{IR}\Pi}$	Increment of the product counter, $S_n^{\mathrm{IR}\Pi}(\sigma, x_0)$	152
D	Minkowski dimension	166

Appendix B

Electric test data

Channel one is yellow (CH1) and measures the driving signal from the function generator. The other is channel three, in magenta (CH3), which measures the voltage of current supply. Above each oscilloscope image is a line of the form "T = Ts, D% duty cycle, I A". Where T and D are the settings for the function generator for cycle period and duty cycle and I is the current read off the power supply display.

B.1 On-chip wires

The current for these tests was supplied by a Delta Elekronika power supply[1] driven in current limited mode via its V_{prog} input (which expects 0–5 V) directly from a RS PRO[2] function generator. A voltage pattern could be set on a the signal generator such that the power supply drove the desired amount of current whenever the signal generator was in the high phase of its duty cycle. The voltage was measured by the oscilloscope at the entrance to the break-out box.

1. Model ES 015-10
2. Model AFG-21005.

- **DC 50**

$\Omega_4 = 7.37\,\Omega$

05: T = 10 s, 10% duty cycle, 1.5 A

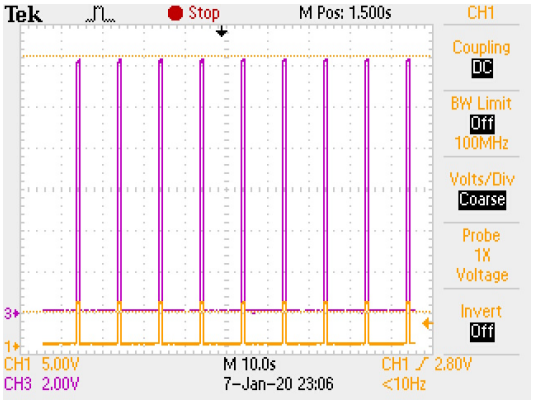

T = 10 s, 30% duty cycle, 0.5 A

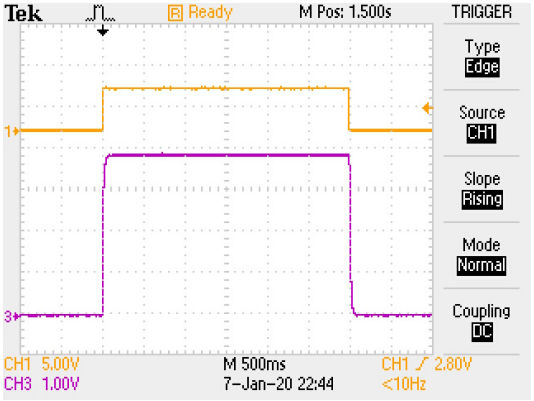

T = 10 s, 30% duty cycle, 1.5 A

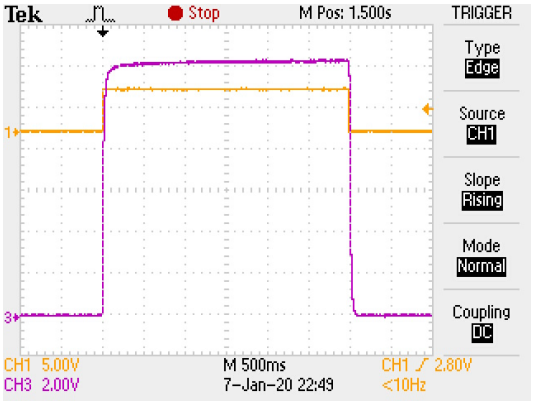

$\Omega_{1.5\,\mathrm{A}} = 8.2\,\Omega$

- **DC 100**

$\Omega_4 = 5.94\,\Omega$

T = 10 s, 10% duty cycle, 1.5 A

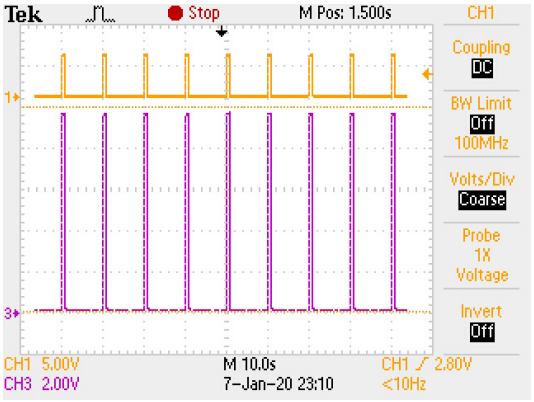

T = 10 s, 30% duty cycle, 0.5 A

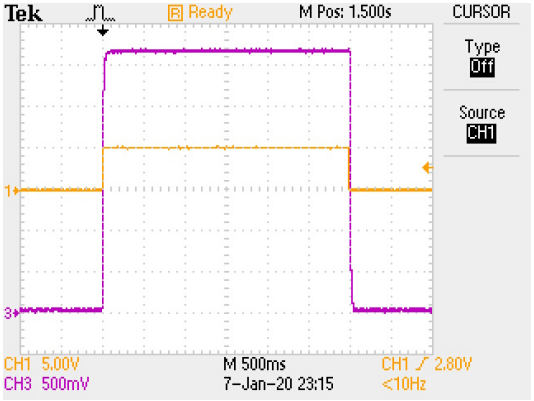

T = 10 s, 30% duty cycle, 1.5 A

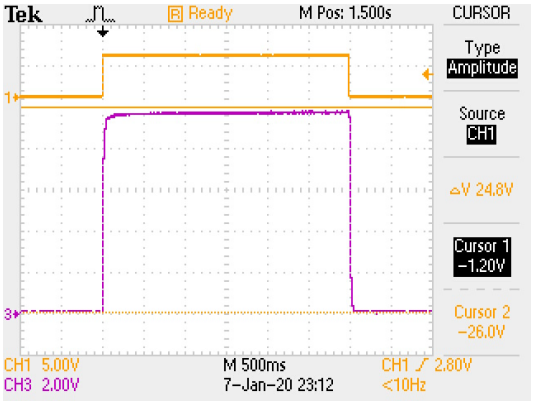

$\Omega_{1.5\,\mathrm{A}} = 6.47\,\Omega$

- **U 1**

$\Omega_4 = 5.88\,\Omega$

T = 10 s, 10% duty cycle, 1.5 A

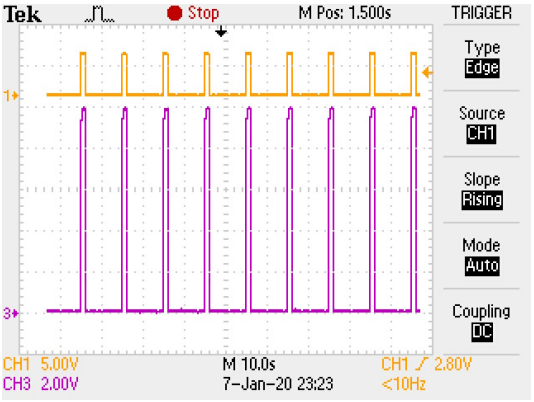

T = 10 s, 30% duty cycle, 0.5 A

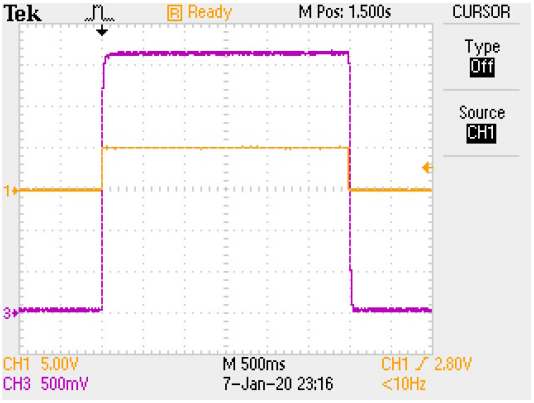

T = 10 s, 30% duty cycle, 1.5 A

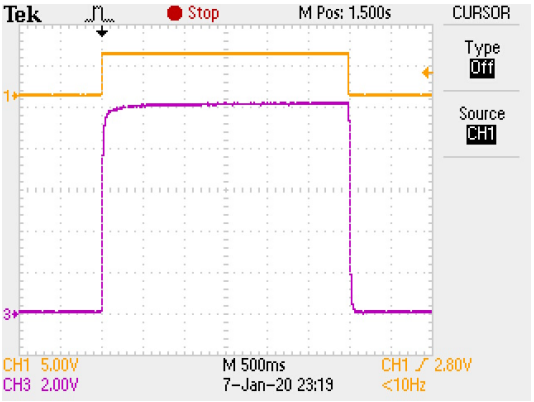

$\Omega_{1.5\,\mathrm{A}} = 6.73\,\Omega$

- **U 2**

$\Omega_4 = 5.66\,\Omega$

T = 10 s, 10% duty cycle, 1.5 A

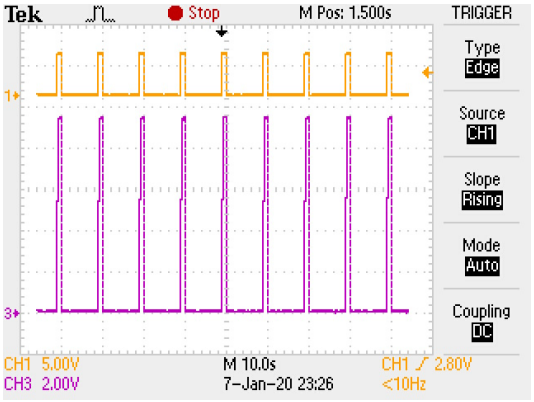

T = 10 s, 30% duty cycle, 0.5 A

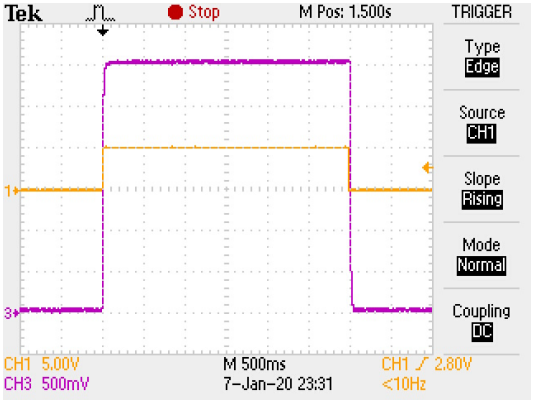

T = 10 s, 30% duty cycle, 1.5 A

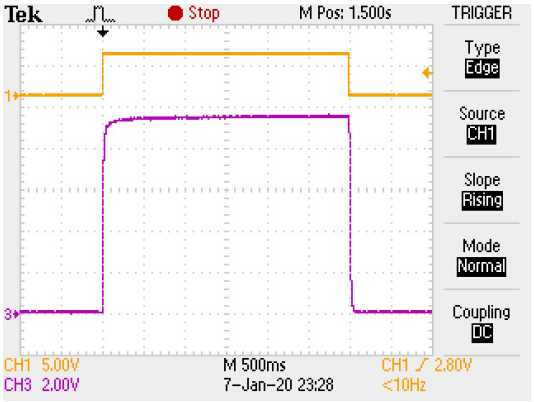

$\Omega_{1.5\,\mathrm{A}} = 6.35\,\Omega$

- **U 3**

$\Omega_4 = 5.59\,\Omega$

T = 10 s, 10% duty cycle, 1.5 A

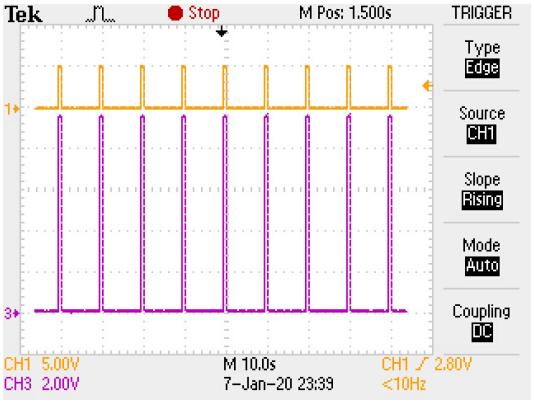

T = 10 s, 30% duty cycle, 0.5 A

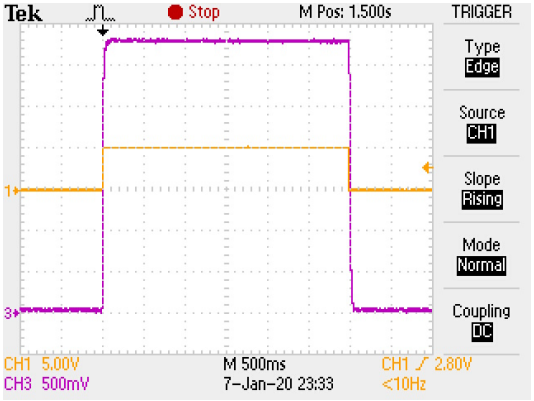

T = 10 s, 30% duty cycle, 1.5 A

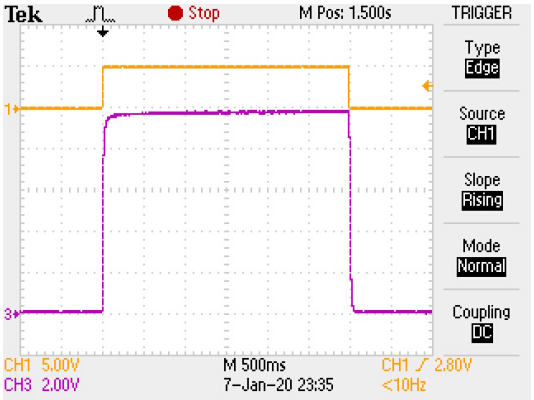

$\Omega_{1.5\,\mathrm{A}} = 6.53\,\Omega$

- **U 4**

$\Omega_4 = 5.33\,\Omega$

T = 10 s, 10% duty cycle, 1.5 A

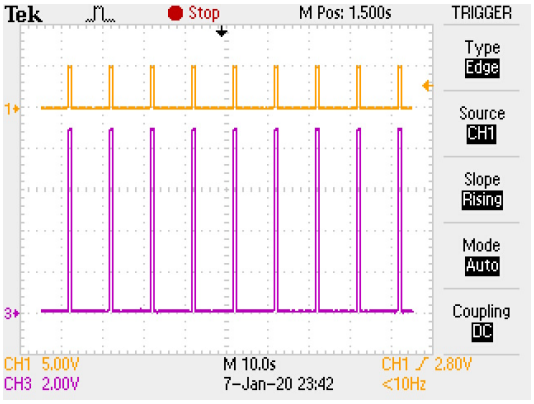

T = 10 s, 30% duty cycle, 0.5 A

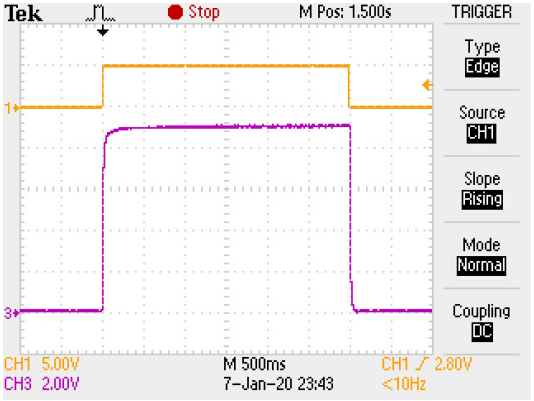

T = 10 s, 30% duty cycle, 1.5 A

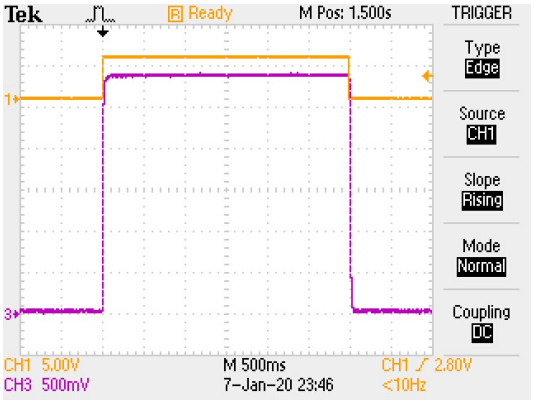

$\Omega_{1.5\,\mathrm{A}} = 6\,\Omega$

- **RF 1**

$\Omega_4 = 18.13\,\Omega$

T = 10 s, 10% duty cycle, 0.5 A

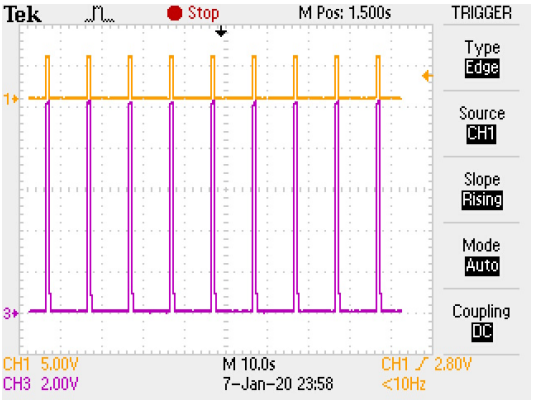

T = 10 s, 30% duty cycle, 0.15 A

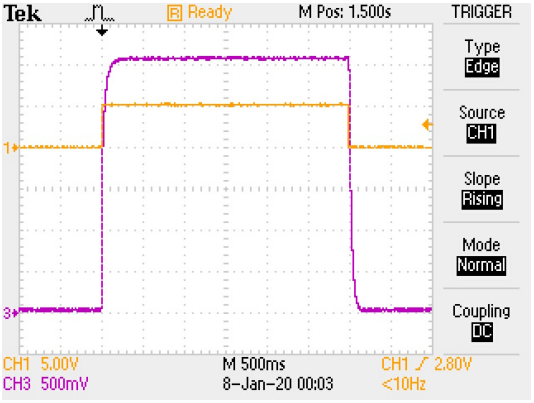

T = 10 s, 30% duty cycle, 0.5 A

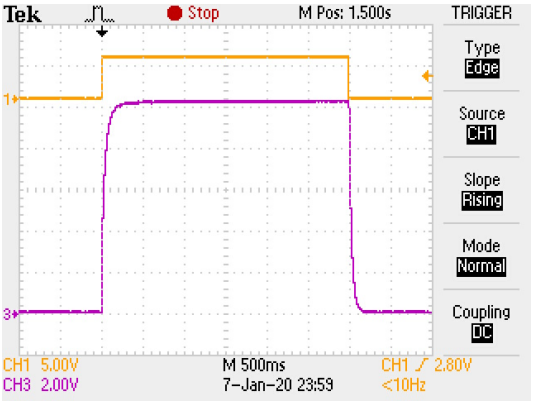

$\Omega_{0.5\,\mathrm{A}} = 20.4\,\Omega$

- **RF 2**

$\Omega_4 = 17.65\,\Omega$

T = 10 s, 10% duty cycle, 0.5 A

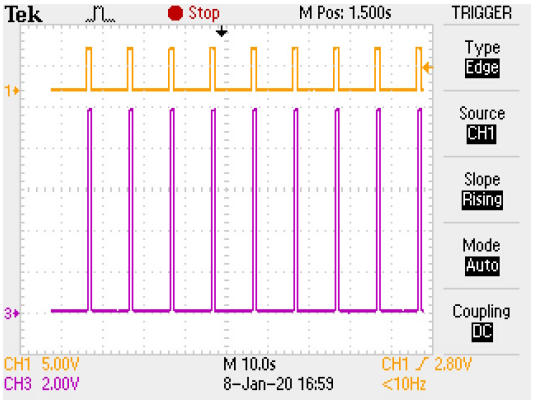

T = 10 s, 30% duty cycle, 0.15 A

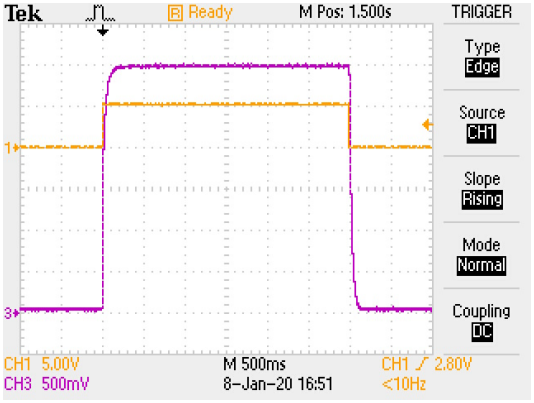

25: T = 10 s, 30% duty cycle, 0.5 A

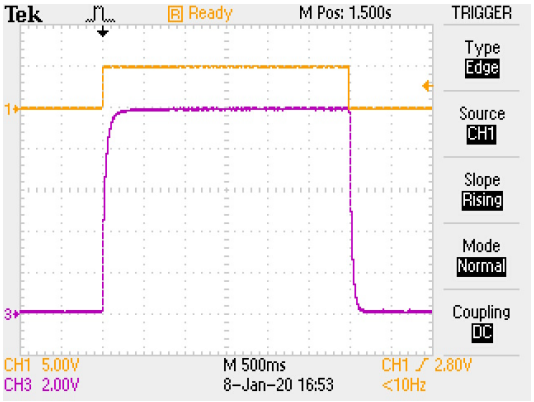

$\Omega_{0.5\,\mathrm{A}} = 19.7\,\Omega$

- **RF 3**

$\Omega_4 = 17.53\,\Omega$

T = 10 s, 10% duty cycle, 0.5 A

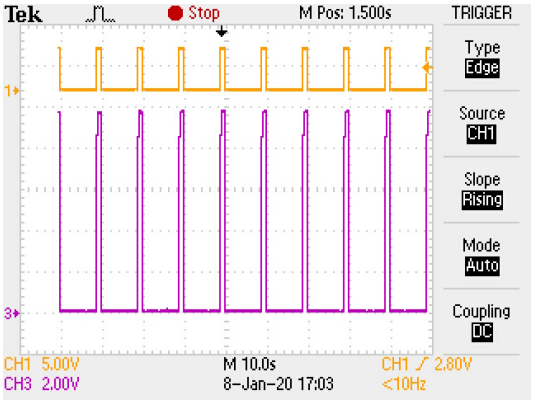

T = 10 s, 30% duty cycle, 0.15 A

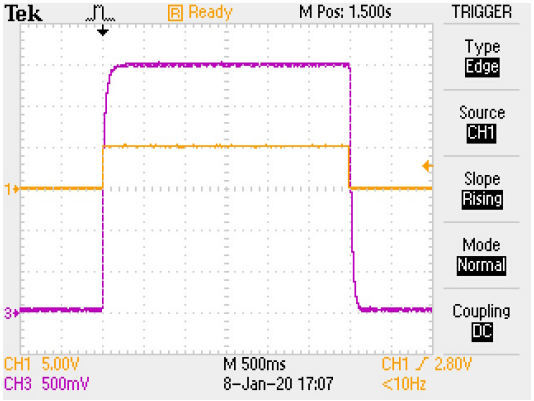

T = 10 s, 30% duty cycle, 0.5 A

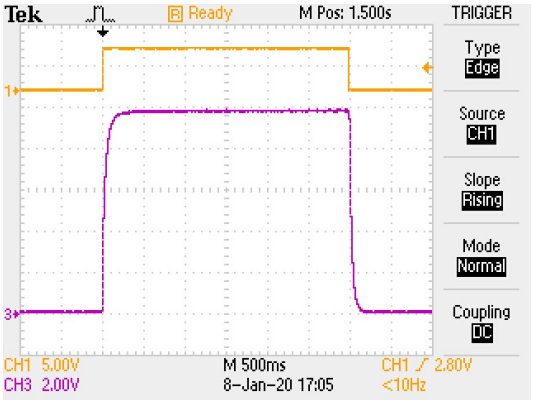

$\Omega_{0.5\,\mathrm{A}} = 19.4\,\Omega$

- **RF 4**

$\Omega_4 = 17.84\,\Omega$

T = 10 s, 10% duty cycle, 0.5 A

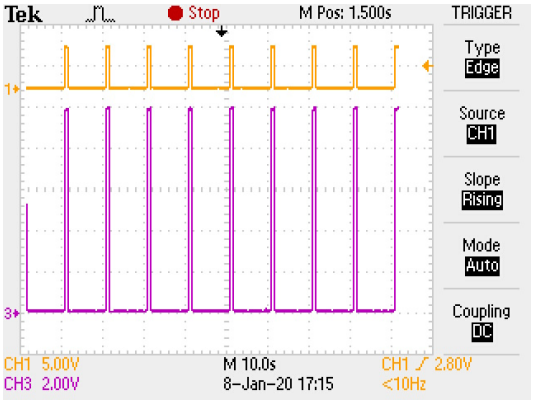

T = 10 s, 30% duty cycle, 0.15 A

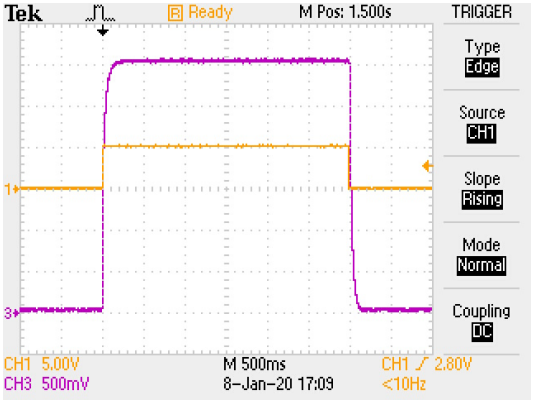

T = 10 s, 30% duty cycle, 0.5 A

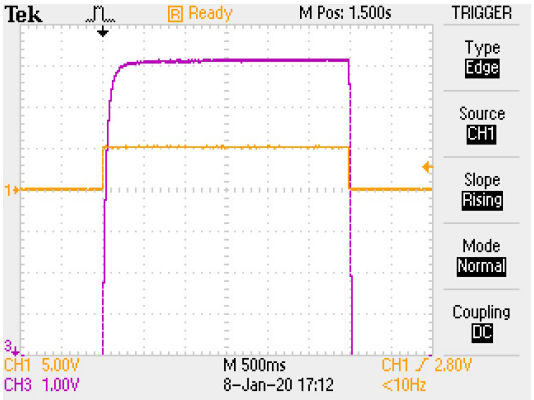

$\Omega_{0.5\,\mathrm{A}} = 20.1\,\Omega$

B.2 Z-wire

The oscilliscope measures the power supply's $V_{monitor}$ voltage through a 1 kΩ, 10 µF RC low-pass filter with a 160 Hz cutoff. Because the 15-100 can supply 0 to 15 V but the $V_{monitor}$ port only has a range of 0 to 5 V, readouts from it must be multiplied by three. In addition to the notation from the last section, these figures include he values Ω_0 and Ω_f.

T = 10 s, 20% duty cycle, 50 A

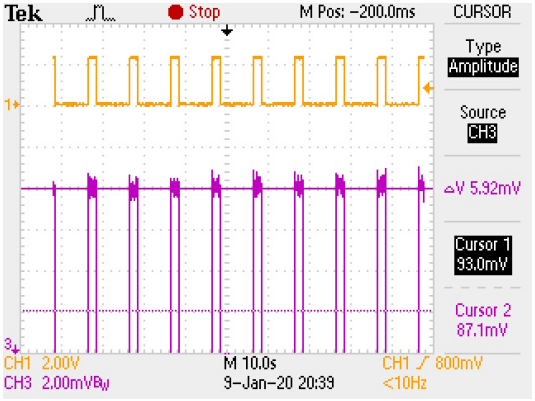

T = 20 s, 20% duty cycle, 25 A

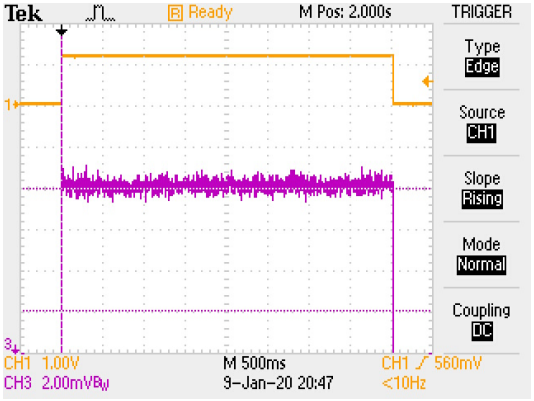

$\Omega_0 = 5.70\,\text{m}\Omega$ $\Omega_f = 5.69\,\text{m}\Omega$

T = 20 s, 20% duty cycle, 50 A

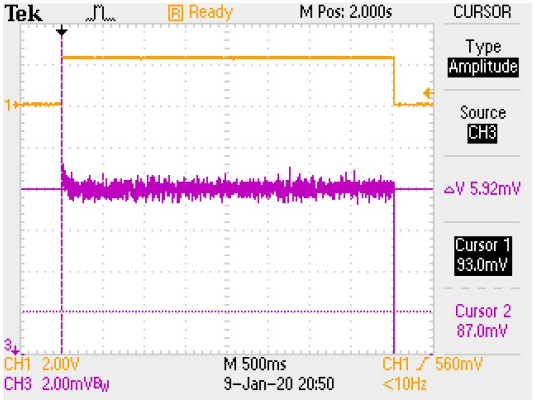

$\Omega_0 = 5.58\,\text{m}\Omega$ $\Omega_f = 5.58\,\text{m}\Omega$

T = 20 s, 20% duty cycle, 75 A

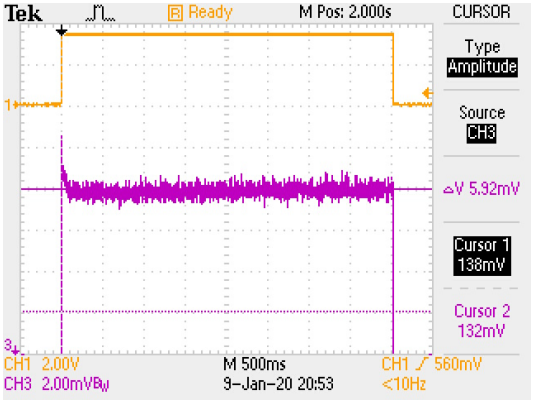

$\Omega_0 = 5.52\,\text{m}\Omega$ $\Omega_f = 5.52\,\text{m}\Omega$

T = 20 s, 20% duty cycle, 100 A

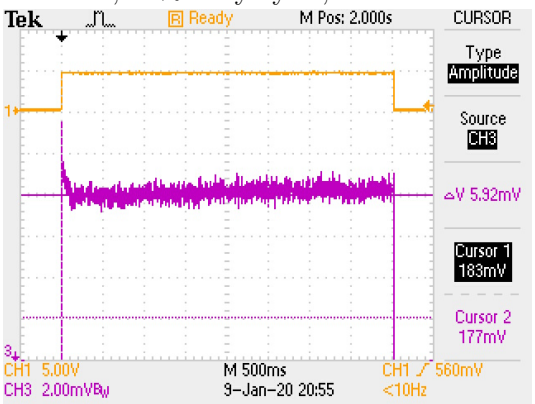

$\Omega_0 = 5.48\,\text{m}\Omega$ $\Omega_f = 5.49\,\text{m}\Omega$

T = 20 s, 20% duty cycle, 112.4 A

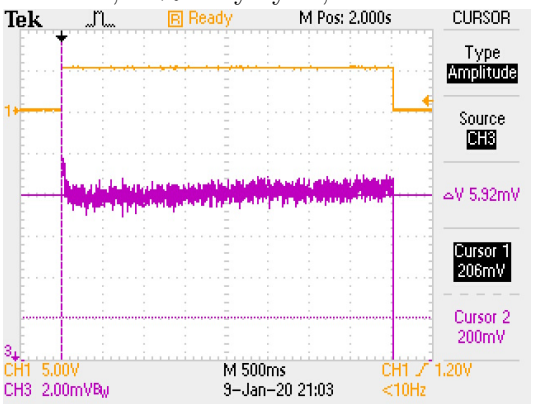

$\Omega_0 = 5.50\,\text{m}\Omega$ $\Omega_f = 5.51\,\text{m}\Omega$

B.3 Evaporation coil

These data were taken with a probe consisting of a single copper wire loop of area of approximately $\pi\,\mathrm{cm}^2$ and an inductance 57 nH. Resistive ballast was supplied by a $50\,\Omega$ load.

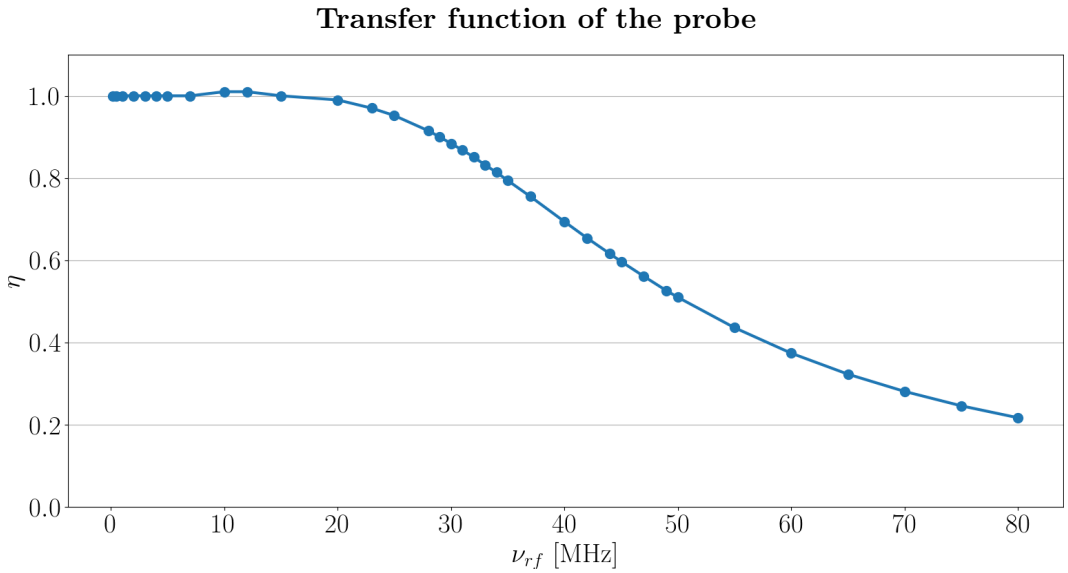

Figure B.1 — Measured gain of the probe used to measure the field strength of the evaporative coil versus driving frequency. Used to correct the raw readings.

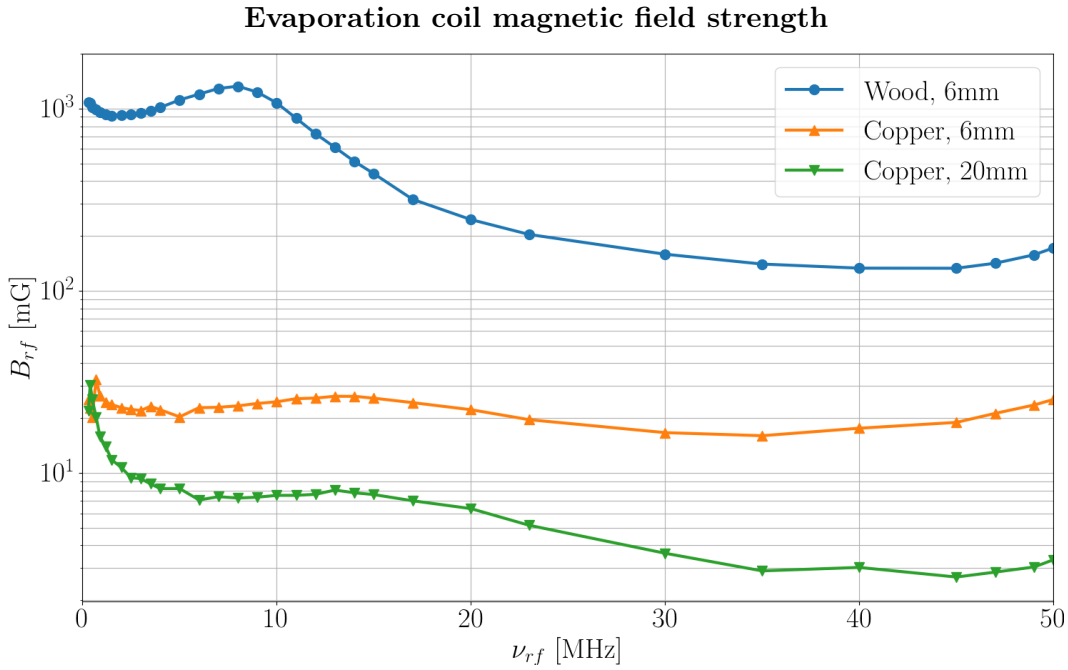

Figure B.2 — Magnetic field generated by the evaporation coil at 6 mm with a wood and copper core and at 20 mm with a copper core, corresponding to our experimental conditions. Coil was run at frequency ν_{rf} with 0.5 W of power.

Appendix C

Objective lens schematics

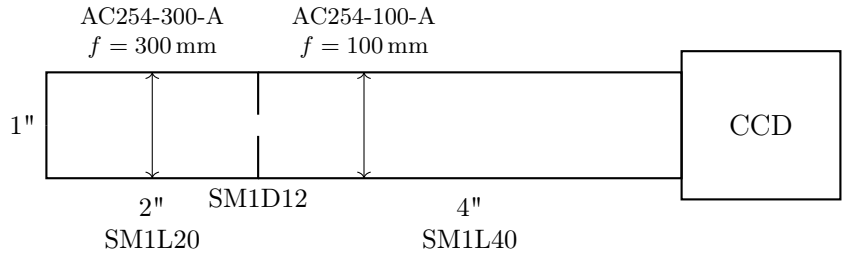

Figure C.1 — Diagram of the low magnification lens. Only tube dimensions to scale. Part numbers from the ThorLabs catalog.

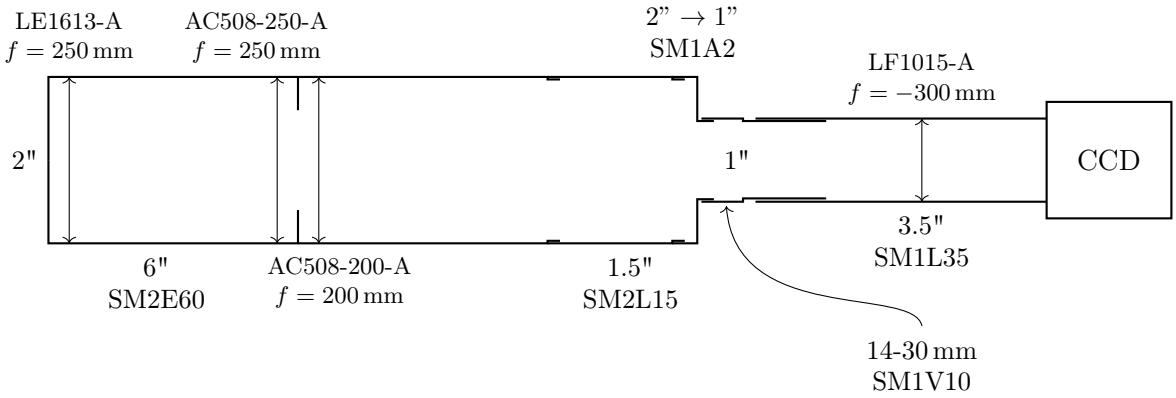

Figure C.2 — Same of the "high" magnification lens for ToF imaging. To convert for in-situ imaging replace the LF1015-A with an LC1582-A, $f = 75\,\mathrm{mm}$, add another SM1L40 between the SM1L35 and the CCD camera and refocus everything according to Table 4.3. The apertures are custom cut blanks and will have to be swapped out.

Appendix D

Supplementary transport profiles

Hyperbolic-tangent profile The idea here is similar to the error function profile: something smooth in all its derivatives that can be tuned between something like a step function and something like the constant velocity profile. It is defined just like the error function profile with one function swapped for the other. It reads

$$y_0(t) = \begin{cases} 0 & \text{for } t \leq 0 \\ \frac{L_1}{2}\left[1 + \tanh\left\{-\gamma\left(\frac{t}{\Delta t}\right)^{-\delta} + \left(1 - \frac{t}{\Delta t}\right)^{-\delta}\right\}\right] & \text{for } t \in]0, \Delta t[\\ L_1 & \text{for } t \geq \Delta t \end{cases} \tag{D.1}$$

with $\gamma = 2^{-3/2-\delta}\sqrt{(\delta+1)(\delta+2)}/\delta$ as before and δ taking the same role as a shape parameter which changes the profile to something resembling a constant velocity one when $\delta \to 0$. This performed very similarly to the error function profile having a clear optimum in its shape parameter around $\delta = .18$, but performing slightly worse than overall.

Sinusoidal acceleration Another way to avoid acceleration spikes while also avoiding spikes in the jerk is to have a profile where the acceleration is sinusoidal. This is written

$$y_0(t) = \begin{cases} 0 & \text{for } t \leq 0 \\ L_1\left[\frac{t}{\Delta t} + (2\pi)^{-1}\sin\left(2\pi\frac{t}{\Delta t}\right)\right] & \text{for } t \in]0, \Delta t[\\ L_1 & \text{for } t \geq \Delta t \end{cases}. \tag{D.2}$$

This did not prove any significant improvement. This suggests that it is not sudden changes in acceleration that cause atom losses or heating, but rather where the atoms are located in the trap as discussed in § 5.3.2 and following.

Constant acceleration with a pause The most creative transport scheme was one where we tried to avoid Majorana losses and heating almost entirely. Imagine the cloud of atoms as liquid sitting at the bottom of a cone-shaped container under gravity then accelerate the container. In the moving frame, acceleration will tilt the container away from the direction of acceleration and the liquid will tend to lie on the lowered side. The strategy is then to stop the trap in the middle of the transport, flinging the atoms across the trap center and then

stop the cloud with the other side of the decelerating potential until trap and atoms come to rest at the end of the transport. Like throwing the liquid from one side of the container to the other and then trying to keep the container moving such that you stop the liquid without making a splash. The profile we attempted to do this with was written

$$y_0(t) = \begin{cases} 0 & \text{for } t \leq 0 \\ 2L_1 \left(\frac{t}{\Delta t - \delta t}\right)^2 & \text{for } t \in \left]0, \frac{\Delta t - \delta t}{2}\right] \\ \frac{L_1}{2} & \text{for } t \in \left]\frac{\Delta t - \delta t}{2}, \frac{\Delta t + \delta t}{2}\right[\\ L_1 \left[1 - 2\left(1 - \frac{t}{\Delta t + \delta t}\right)^2\right] & \text{for } t \in \left]\frac{\Delta t + \delta t}{2}, \Delta t + \delta t\right[\\ L_1 & \text{for } t \geq \Delta t \end{cases} . \tag{D.3}$$

This is just the constant acceleration profile defined in equation (5.13) with a pause of length δt in the middle at $L_1/2, \Delta t/2$. Several attempts at scanning Δt and δt to find the right combination all lead to a total loss of the atom cloud, even at fairly slow transport times on the order of a second.

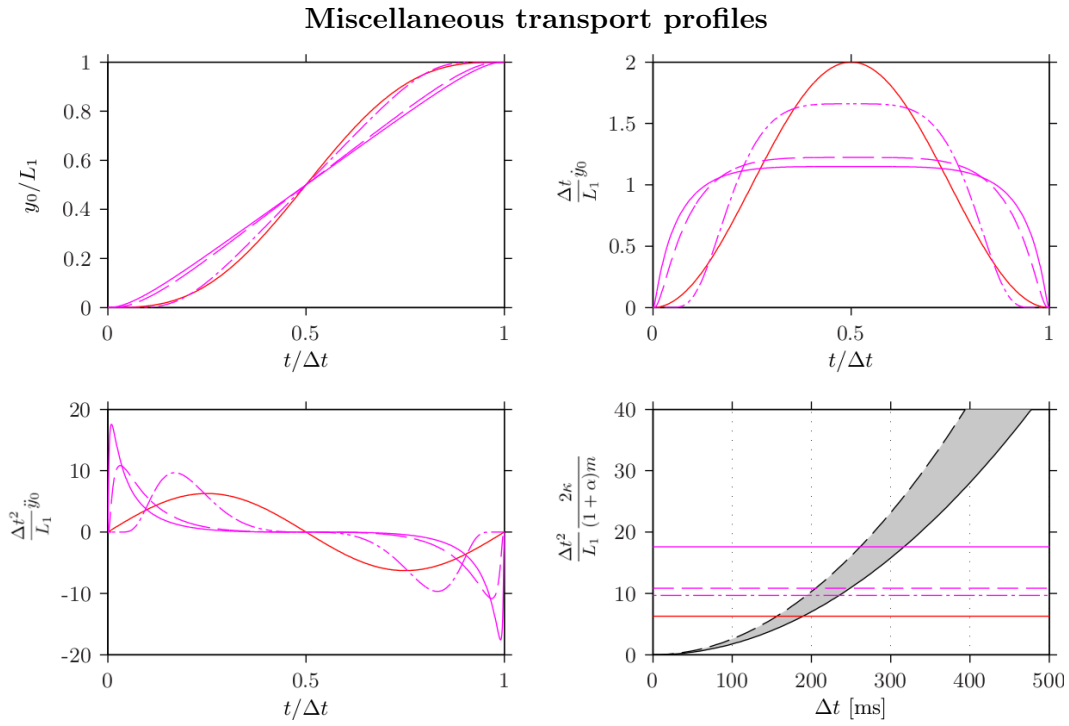

Figure D.1 — Same as Figure 5.5 but for the sinusoidal-acceleration (red line) and hyperbolic-tangent (magenta lines) profiles. The line style denotes different values of the shape parameter, $\delta = 0.2$ (solid), $\delta = 0.3$ (dashed), and $\delta = 0.9$ (dashed-dotted).

Appendix E

One-dimensional time-of-flight density profile

Here I describe the calculation for the time-dependent one-dimensional density profile of a cloud released from the magnetic trap mentioned in § 5.2.1.2 and give the exact, analytic formula.

To construct a ToF profile, we convolve the velocity profile with the density in the trap (equation (5.3)). This amounts to averaging the velocity distribution over the volume of the trap, weighted by the density distribution in the trap which serves as a set of initial positions for the atoms.

When integrated along the imaging axis, y, it behaves like a Gaussian because y only appears in $\boldsymbol{r}$. This gives a factor of $\sqrt{2\pi r_t^2}$ and a distribution function for the image of

$$\begin{aligned} n(x,z,t) &= \int_{-\infty}^{\infty} \mathrm{d}y \, \left(2\pi r_t^2\right)^{-3/2} \int \mathrm{d}\boldsymbol{u} \, n(\boldsymbol{u}) \exp\left(-\frac{(\boldsymbol{r}-\boldsymbol{u})^2}{2r_t^2}\right) \\ &= \frac{C}{2\pi r_t^2} \int \mathrm{d}\boldsymbol{u} \, e^{-u/R} e^{-(u_x-x)^2/2r_t^2} e^{-(u_z-z)^2/2r_t^2} \end{aligned} \tag{E.1}$$

where $C = N/4\pi R^3$ is the normalization factor from the distribution in the trap, $R \propto k_B T$ is the radius of mean density as defined in equation (5.4).

To get a one-dimensional profile that can be fit to integrated profile from an experimental image to extract a temperature, one more integration along either x or z must be performed. Then the integral over $\boldsymbol{u}$ must be done explicitly. The trick of switching to spherical coordinates to avoid improper integration limits as was done in arriving at equation (5.6) can be used here. Integrating equation (E.1) along z results in

$$\begin{aligned} n(x,t) = \frac{r_t}{2\sqrt{2\pi}R^2} e^{-x^2/2r_t^2} + \frac{1}{8R^3} e^{r_t^2/2R^2} \times \ldots \\ \left[\left(R^2 + Rx - r_t^2\right) \operatorname{erfc}\left(\frac{1}{\sqrt{2}}\left[\frac{r_t}{R} - \frac{x}{r_t}\right]\right) e^{-x/R} + \left(R^2 - Rx - r_t^2\right) \operatorname{erfc}\left(\frac{1}{\sqrt{2}}\left[\frac{r_t}{R} + \frac{x}{r_t}\right]\right) e^{x/R}\right]. \end{aligned} \tag{E.2}$$

The vertical axis profile differs from the horizontal one in that R is replaced by $R/2$ due to the steeper trap gradient in that direction. These can both be fit to a doubly integrated profile and the temperature extracted from r_t and R.

Bibliography

[1] Ben Ali, D. *Design and construction of a cold atoms experiment: towards sodium condensates on a chip.* Ph.D. thesis, Université Paris 13, Sorbonne Paris Cité, Paris (May 2016).

[2] Greiner, M., Bloch, I., Hänsch, T., and Esslinger, T. "Magnetic transport of trapped cold atoms over a large distance." *Physical Review A* (2001). **63**, pp. 031401, 1–4.

[3] Silva, C. E. *Invitation to Ergodic Theory.* Number 42 in Student Mathematical Library, American Mathematical Society, Providence, R.I. (2008).

[4] Brin, M. and Stuck, G. *Introduction to Dynamical Systems.* Cambridge University Press, Cambridge, U.K. (2002).

[5] Einstein, A. "Quantentheorie des einatomigen idealen Gases." *Sitzungsberichte der Preussischen Akademie der Wissenschaften zu Berlin. Physikalisch-mathematische Klasse* (July 1924). **22**, pp. 261–267.

[6] Einstein, A. "Quantentheorie des einatomigen idealen Gases. Zweite Abhandlung." *Sitzungsberichte der Preussischen Akademie der Wissenschaften zu Berlin. Physikalisch-mathematische Klasse* (January 1925). **1**(3), pp. 3–14.

[7] Bose, S. N. "Plancks Gesetz und Lichtquantenhypothese." *Zeitschrift für Physik* (December 1924). **26**(1), pp. 178–181.

[8] Pauli, W. "Über den Zusammenhang des Abschlusses der Elektronengruppen im Atom mit der Komplexstruktur der Spektren." *Zeitschrift für Physik* (February 1925). **31**(1), pp. 765–783.

[9] London, F. "On the Bose–Einstein Condensation." *Physical Review* (December 1938). **54**(11), pp. 947–954.

[10] Phillips, W. D. and Metcalf, H. "Laser deceleration of an atomic beam." *Physical Review Letters* (March 1982). **48**(9), pp. 596–599.

[11] Chu, S., Hollberg, L., Bjorkholm, J. E., Cable, A., and Ashkin, A. "Three-dimensional viscous confinement and cooling of atoms by resonance radiation pressure." *Physical Review Letters* (1985). **55**, pp. 48–51.

[12] Raab, E. L., Prentiss, M., Cable, A., Chu, S., and Pritchard, D. E. "Trapping of Neutral Sodium Atoms with Radiation Pressure." *Physical Review Letters* (December 1987). **59**(23), pp. 2631–2634.

[13] Dalibard, J. and Cohen-Tannoudji, C. "Laser cooling below the Doppler limit by polarization gradients: simple theoretical models." *Journal of the Optical Society of America B* (November 1989). **6**(11), p. 2023.

[14] Chu, S. "Nobel Lecture: The manipulation of neutral particles." *Reviews of Modern Physics* (July 1998). **70**(3), pp. 685–706.

[15] Cohen-Tannoudji, C. N. "Nobel Lecture: Manipulating atoms with photons." *Reviews of Modern Physics* (July 1998). **70**(3), pp. 707–719.

[16] Phillips, W. D. "Nobel Lecture: Laser cooling and trapping of neutral atoms." *Reviews of Modern Physics* (July 1998). **70**(3), pp. 721–741.

[17] Davis, K. B., Mewes, M.-O., Joffe, M. A., Andrews, M. R., and Ketterle, W. "Evaporative Cooling of Sodium Atoms." *Physical Review Letters* (June 1995). **74**(26), pp. 5202–5205.

[18] Migdall, A. L., Prodan, J. V., Phillips, W. D., Bergeman, T. H., and Metcalf, H. J. "First Observation of Magnetically Trapped Neutral Atoms." *Physical Review Letters* (June 1985). **54**(24), pp. 2596–2599.

[19] Anderson, M. H., Ensher, J. R., Matthews, M. R., Wieman, C. E., and Cornell, E. A. "Observation of Bose-Einstein condensation in a dilute atomic vapor." *Science* (July 1995). **269**(5221), pp. 198–201.

[20] Cornell, E. A. and Wieman, C. E. "Nobel Lecture: Bose-Einstein condensation in a dilute gas, the first 70 years and some recent experiments." *Reviews of Modern Physics* (August 2002). **74**(3), pp. 875–893.

[21] Ketterle, W. "Nobel lecture: When atoms behave as waves: Bose-Einstein condensation and the atom laser." *Reviews of Modern Physics* (November 2002). **74**(4), pp. 1131–1151.

[22] Weber, T., Herbig, J., Mark, M., Nägerl, H.-C., and Grimm, R. "Bose–Einstein Condensation of Cesium." *Science* (2003). **299**(5604), pp. 232–235.

[23] Jochim, S. "Bose–Einstein Condensation of Molecules." *Science* (December 2003). **302**(5653), pp. 2101–2103.

[24] Misochko, O., Hase, M., Ishioka, K., and Kitajima, M. "Transient Bose–Einstein condensation of phonons." *Physics Letters A* (February 2004). **321**(5-6), pp. 381–387.

[25] Oosawa, A., Ishii, M., and Tanaka, H. "Field-induced three-dimensional magnetic ordering in the spin-gap system." *Journal of Physics: Condensed Matter* (January 1999). **11**(1), pp. 265–271.

[26] Klaers, J., Schmitt, J., Vewinger, F., and Weitz, M. "Bose–Einstein condensation of photons in an optical microcavity." *Nature* (November 2010). **468**(7323), pp. 545–548.

[27] Kasprzak, J., Richard, M., Kundermann, S., Baas, A., Jeambrun, P., Keeling, J. M. J., Marchetti, F. M., Szymańska, M. H., André, R., Staehli, J. L., Savona, V., Littlewood, P. B., Deveaud, B., and Dang, L. S. "Bose–Einstein condensation of exciton polaritons." *Nature* (September 2006). **443**(7110), pp. 409–414.

[28] Deng, H., Haug, H., and Yamamoto, Y. "Exciton-polariton Bose-Einstein condensation." *Reviews of Modern Physics* (May 2010). **82**(2), pp. 1489–1537.

[29] Byrnes, T., Rosseau, D., Khosla, M., Pyrkov, A., Thomasen, A., Mukai, T., Koyama, S., Abdelrahman, A., and Ilo-Okeke, E. "Macroscopic quantum information processing using spin coherent states." *Optics Communications* (February 2015). **337**, pp. 102–109.

[30] Eckel, S., Kumar, A., Jacobson, T., Spielman, I., and Campbell, G. "A Rapidly Expanding Bose-Einstein Condensate: An Expanding Universe in the Lab." *Physical Review X* (April 2018). **8**(2), p. 021021.

[31] Ketterle, W., Durfee, D. S., and Stamper-Kurn, D. M. "Making, probing and understanding Bose-Einstein condensates." In *Proceedings of the International School of Physics "Enrico Fermi" Course CXL*, edited by M. Inguscio, S. Stringari, and C. Wieman. IOS Press, Varenna, Lake Como (1998) pp. 67–176. ArXiv: cond-mat/9904034.

[32] Chin, C., Grimm, R., Julienne, P., and Tiesinga, E. "Feshbach Resonances in Ultracold Gases." *Reviews of Modern Physics* (April 2010). **82**(2), pp. 1225–1286. ArXiv: 0812.1496.

[33] Rice, O. K. "Predissociation and the Crossing of Molecular Potential Energy Curves." *The Journal of Chemical Physics* (June 1933). **1**(6), pp. 375–389.

[34] Feshbach, H. "Unified theory of nuclear reactions." *Annals of Physics* (December 1958). **5**(4), pp. 357–390.

[35] Fano, U. "Effects of Configuration Interaction on Intensities and Phase Shifts." *Physical Review* (December 1961). **124**(6), pp. 1866–1878.

[36] Stwalley, W. C. "Stability of Spin-Aligned Hydrogen at Low Temperatures and High Magnetic Fields: New Field-Dependent Scattering Resonances and Predissociations." *Physical Review Letters* (December 1976). **37**(24), pp. 1628–1631.

[37] Tiesinga, E., Verhaar, B. J., and Stoof, H. T. C. "Threshold and resonance phenomena in ultracold ground-state collisions." *Physical Review A* (May 1993). **47**(5), pp. 4114–4122.

[38] Moerdijk, A. J., Verhaar, B. J., and Axelsson, A. "Resonances in ultracold collisions of ^{6}Li, ^{7}Li, and ^{23}Na." *Physical Review A* (June 1995). **51**(6), pp. 4852–4861.

[39] Courteille, P., Freeland, R. S., Heinzen, D. J., van Abeelen, F. A., and Verhaar, B. J. "Observation of a Feshbach Resonance in Cold Atom Scattering." *Physical Review Letters* (July 1998). **81**(1), pp. 69–72.

[40] Shkedrov, C., Florshaim, Y., Ness, G., Gandman, A., and Sagi, Y. "High-Sensitivity rf Spectroscopy of a Strongly Interacting Fermi Gas." *Physical Review Letters* (August 2018). **121**(9), p. 093402.

[41] Fedichev, P. O., Kagan, Y., Shlyapnikov, G. V., and Walraven, J. T. M. "Influence of Nearly Resonant Light on the Scattering Length in Low-Temperature Atomic Gases." *Physical Review Letters* (September 1996). **77**(14), pp. 2913–2916.

[42] Jones, K. M., Tiesinga, E., Lett, P. D., and Julienne, P. S. "Ultracold photoassociation spectroscopy: Long-range molecules and atomic scattering." *Reviews of Modern Physics* (May 2006). **78**(2), pp. 483–535.

[43] Papoular, D. J., Shlyapnikov, G. V., and Dalibard, J. "Microwave-induced Fano-Feshbach resonances." *Physical Review A* (April 2010). **81**(4), p. 041603.

[44] Folman, R., Krüger, P., Cassettari, D., Hessmo, B., Maier, T., and Schmiedmayer, J. "Controlling Cold Atoms using Nanofabricated Surfaces: Atom Chips." *Physical Review Letters* (May 2000). **84**(20), pp. 4749–4752.

[45] Reichel, J. and Vuletić, V. (Editors). *Atom Chips*. Wiley-VCH-Verl, Weinheim (2011).

[46] Simons, R. N. *Coplanar Waveguide Circuits, Components, and Systems*. John Wiley & Sons, Incorporated., Hoboken, N.J. (2002).

[47] Bouchoule, I., Westbrook, C. I., and van Druten, N. J. "Atom Chips and One-Dimensional Bose Gases." In *Atom Chips*, edited by J. Reichel and V. Vuletić, pp. 331–363. Wiley-VCH Verlag GmbH & Co. KGaA, Weinheim, Germany, first edition (February 2011).

[48] Lieb, E. H. and Liniger, W. "Exact Analysis of an Interacting Bose Gas. I. The General Solution and the Ground State." *Physical Review* (May 1963). **130**(4), pp. 1605–1616.

[49] Lieb, E. H. "Exact Analysis of an Interacting Bose Gas. II. The Excitation Spectrum." *Physical Review* (May 1963). **130**(4), pp. 1616–1624.

[50] Yang, C. N. and Yang, C. P. "Thermodynamics of a One-Dimensional System of Bosons with Repulsive Delta-Function Interaction." *Journal of Mathematical Physics* (July 1969). **10**(7), pp. 1115–1122.

[51] Paredes, B., Widera, A., Murg, V., Mandel, O., Fölling, S., Cirac, I., Shlyapnikov, G. V., Hänsch, T. W., and Bloch, I. "Tonks-Girardeau gas of ultracold atoms in an optical lattice." *Nature* (May 2004). **429**(6989), pp. 277–281.

[52] Kinoshita, T., Wenger, T., and Weiss, D. S. "Observation of a one-dimensional Tonks-Girardeau gas." *Science* (2004). **305**(5687), pp. 1125–1128.

[53] Boéris, G., Gori, L., Hoogerland, M. D., Kumar, A., Lucioni, E., Tanzi, L., Inguscio, M., Giamarchi, T., D'Errico, C., Carleo, G., Modugno, G., and Sanchez-Palencia, L. "Mott transition for strongly interacting one-dimensional bosons in a shallow periodic potential." *Physical Review A* (January 2016). **93**(1), p. 011601.

[54] Jacqmin, T., Armijo, J., Berrada, T., Kheruntsyan, K. V., and Bouchoule, I. "Sub-Poissonian Fluctuations in a 1D Bose Gas: From the Quantum Quasicondensate to the Strongly Interacting Regime." *Physical Review Letters* (June 2011). **106**(23), p. 230405.

[55] Bouchoule, I., Trebbia, J.-B., and Garrido Alzar, C. L. "Limitations of the modulation method to smooth wire-guide roughness." *Physical Review A* (February 2008). **77**(2), p. 023624.

[56] Kheruntsyan, K. V., Gangardt, D. M., Drummond, P. D., and Shlyapnikov, G. V. "Pair Correlations in a Finite-Temperature 1D Bose Gas." *Physical Review Letters* (July 2003). **91**(4), pp. 040403, 1–4.

[57] Haldane, F. D. M. "Effective Harmonic-Fluid Approach to Low-Energy Properties of One-Dimensional Quantum Fluids." *Physical Review Letters* (December 1981). **47**(25), pp. 1840–1843.

[58] von Plato, J. "Boltzmann's ergodic hypothesis." *Archive for History of Exact Sciences* (Mar 1991). **42**(1), pp. 71–89.

[59] Boltzmann, L. *Wissenschaftliche Abhandlungen*, volume 2. Cambridge University Press, Cambridge, UK (2012).

[60] Plancherel, M. "Beweis der Unmöglichkeit ergodischer mechanischer Systeme." *Annalen der Physik* (1913). **347**(15), pp. 1061–1063.

[61] Rosenthal, A. "Beweis der Unmöglichkeit ergodischer Gassysteme." *Annalen der Physik* (1913). **347**(14), pp. 796–806.

[62] Ehrenfest, P. and Ehrenfest, T. *The Conceptual Foundations of the Statistical Approach in Mechanics*. Dover Books on Physics and Chemistry, Dover Publications, New York, N.Y. (1990).

[63] Neumann, J. v. "Proof of the quasi-ergodic hypothesis." *Proceedings of the National Academy of Sciences* (1932). **18**(1), pp. 70–82.

[64] Birkhoff, G. D. "Proof of the ergodic theorem." *Proceedings of the National Academy of Sciences* (Dec 1931). **17**(12), pp. 656–660.

[65] Alicki, R. and Fannes, M. *Quantum Dynamical Systems*. Oxford University Press, Oxford, UK (2001).

[66] Tabachnikov, S. *Geometry and Billiards*. Number 30 in Student Mathematical Library, American Mathematical Society, Providence, R.I. (2005).

[67] Coxeter, H. S. M. "Discrete groups generated by reflections." *The Annals of Mathematics* (July 1934). **35**(3), pp. 588–621.

[68] Gaudin, M. *La fonction d'onde de Bethe*. Masson, Paris (1983).

[69] Harshman, N. L., Olshanii, M., Dehkharghani, A. S., Volosniev, A. G., Jackson, S. G., and Zinner, N. T. "Integrable Families of Hard-Core Particles with Unequal Masses in a One-Dimensional Harmonic Trap." *Physical Review X* (October 2017). **7**(4), p. 041001.

[70] Wang, J., Casati, G., and Prosen, T. "Nonergodicity and localization of invariant measure for two colliding masses." *Physical Review E* (2014). **89**(4), p. 042918.

[71] Landau, L. D. and Lifshitz, E. M. *Quantum Mechanics: Non-relativistic Theory*, volume 3 of *Course of Theoretical Physics*. Pergamon Press, Oxford, UK, third edition (1977).

[72] Landau, L. D. and Lifshitz, E. M. *Statistical Physics: Part 1*, volume 5 of *Course of Theoretical Physics*. Pergamon Press, Oxford, UK, third edition (1980).

[73] Pitaevskii, L. and Stringari, S. *Bose–Einstein Condensation*. University Press, Oxford, UK, first edition (2003).

[74] Perrin, H. "Low-dimensional Bose gases Part 1: BEC and interactions." Course slides, available online at `http://www-lpl.univ-paris13.fr/bec/BEC/Teaching/lecture1LowD.pdf` (2009).

[75] Metcalf, H. J. and Van der Straten, P. *Laser Cooling and Trapping*. Graduate Texts in Contemporary Physics, Springer, New York, N.Y., first edition (1999).

[76] Bloch, I., Dalibard, J., and Zwerger, W. "Many-body physics with ultracold gases." *Reviews of Modern Physics* (July 2008). **80**(3), pp. 885–964.

[77] Petrov, D., Gangardt, D., and Shlyapnikov, G. "Low-dimensional trapped gases." In *Proceedings of the Euroschool on quantum gases in low dimensions, Les Houches 2003*, volume 116, edited by L. Pricoupenko, H. Perrin, and M. Olshanii. J. Phys. IV (2004) p. 5.

[78] Abramowitz, M. and Stegun, I. A. (Editors). *Handbook of Mathematical Functions: with Formulas, Graphs, and Mathematical Tables*. Dover Books on Mathematics, Dover Publications, New York, N.Y., 9th edition (2013).

[79] Lewin, L. *Polylogarithms and Associated Functions*. North Holland, New York, N.Y. (1981).

[80] Sutherland, B. *Beautiful Models: 70 Years of Exactly Solved Quantum Many-Body Problems*. World Scientific, Singapore (2004).

[81] Ferrenberg, A. M., Xu, J., and Landau, D. P. "Pushing the limits of Monte Carlo simulations for the three-dimensional Ising model." *Physical Review E* (April 2018). **97**(4), p. 043301.

[82] Gross, E. P. "Structure of a quantized vortex in boson systems." *Il Nuovo Cimento* (May 1961). **20**(3), pp. 454–477.

[83] Pitaevskii, L. P. "Vortex lines in an imperfect Bose gas." *Soviet Physics, Journal of Experimentam and Theoretical Physics* (1961). **13**(2), pp. 451–454.

[84] Bogoliubov, N. "On the theory of superfluidity." *Journal of Physics of the USSR* (1947). **XI**(1), pp. 23–32.

[85] Görlitz, A., Vogels, J. M., Leanhardt, A. E., Raman, C., Gustavson, T. L., Abo-Shaeer, J. R., Chikkatur, A. P., Gupta, S., Inouye, S., Rosenband, T., and Ketterle, W. "Realization of Bose-Einstein Condensates in Lower Dimensions." *Physical Review Letters* (September 2001). **87**(13), p. 130402.

[86] Girardeau, M. "Relationship between systems of impenetrable bosons and fermions in one dimension." *Journal of Mathematical Physics* (November 1960). **1**(6), pp. 516–523.

[87] McGuire, J. B. "Study of exactly soluble one-dimensional n-body problems." *Journal of Mathematical Physics* (May 1964). **5**(5), pp. 622–636.

[88] Dettmer, S., Hellweg, D., Ryytty, P., Arlt, J. J., Ertmer, W., Sengstock, K., Petrov, D. S., Shlyapnikov, G. V., Kreutzmann, H., Santos, L., and Lewenstein, M. "Observation of Phase Fluctuations in Elongated Bose-Einstein Condensates." *Physical Review Letters* (October 2001). **87**(16), p. 160406.

[89] Gerbier, F. "Quasi-1D Bose-Einstein condensates in the dimensional crossover regime." *(EPL) Europhysics Letters* (June 2004). **66**(6), pp. 771–777.

[90] Gerbier, F., Thywissen, J. H., Richard, S., Hugbart, M., Bouyer, P., and Aspect, A. "Critical Temperature of a Trapped, Weakly Interacting Bose Gas." *Physical Review Letters* (January 2004). **92**(3), p. 030405.

[91] van Amerongen, A. H., van Es, J. J. P., Wicke, P., Kheruntsyan, K. V., and van Druten, N. J. "Yang-Yang Thermodynamics on an Atom Chip." *Physical Review Letters* (March 2008). **100**(9), p. 090402.

[92] Estève, J., Trebbia, J.-B., Schumm, T., Aspect, A., Westbrook, C. I., and Bouchoule, I. "Observations of Density Fluctuations in an Elongated Bose Gas: Ideal Gas and Quasicondensate Regimes." *Physical Review Letters* (April 2006). **96**(13), p. 130403.

[93] Olshanii, M. "Atomic Scattering in the Presence of an External Confinement and a Gas of Impenetrable Bosons." *Physical Review Letters* (August 1998). **81**(5), pp. 938–941.

[94] Trebbia, J.-B., Esteve, J., Westbrook, C. I., and Bouchoule, I. "Experimental Evidence for the Breakdown of a Hartree-Fock Approach in a Weakly Interacting Bose Gas." *Physical Review Letters* (December 2006). **97**(25), p. 250403.

[95] Gaudin, M. "Boundary Energy of a Bose Gas in One Dimension." *Physical Review A* (July 1971). **4**(1), pp. 386–394.

[96] van Diejen, J. F. and Vinet, L. *Calogero–Moser–Sutherland Models.* CRM Series in Mathematical Physics, Springer Science & Business Media, New York, N.Y. (2012).

[97] Bethe, H. "Zur Theorie der Metalle: I. Eigenwerte und Eigenfunktionen der linearen Atomkette." *Zeitschrift für Physik* (March 1931). **71**(3-4), pp. 205–226.

[98] Sutherland, B. "Nondiffractive scattering: Scattering from kaleidoscopes." *Journal of Mathematical Physics* (July 1980). **21**(7), pp. 1770–1775.

[99] Olshanii, M. and Jackson, S. G. "An exactly solvable quantum four-body problem associated with the symmetries of an octacube." *New Journal of Physics* (October 2015). **17**(10), p. 105005.

[100] Pitaevskii, L. and Stringari, S. *Bose–Einstein Condensation and Superfluidity.* Oxford University Press, Oxford, UK (January 2016).

[101] Ristivojevic, Z. "Excitation Spectrum of the Lieb-Liniger Model." *Physical Review Letters* (July 2014). **113**(1), p. 015301.

[102] Astrakharchik, G. E. and Pitaevskii, L. P. "Lieb's soliton-like excitations in harmonic traps." *EPL (Europhysics Letters)* (May 2013). **102**(3), p. 30004.

[103] Castro-Alvaredo, O. A., Doyon, B., and Yoshimura, T. "Emergent Hydrodynamics in Integrable Quantum Systems Out of Equilibrium." *Physical Review X* (December 2016). **6**(4), p. 041065.

[104] Bertini, B., Collura, M., De Nardis, J., and Fagotti, M. "Transport in Out-of-Equilibrium X X Z Chains: Exact Profiles of Charges and Currents." *Physical Review Letters* (November 2016). **117**(20), p. 207201.

[105] Schemmer, M., Bouchoule, I., Doyon, B., and Dubail, J. "Generalized Hydrodynamics on an Atom Chip." *Physical Review Letters* (March 2019). **122**(9), p. 090601.

[106] Loudon, R. *The Quantum Theory of Light.* Oxford Science Publications, Oxford University Press, Oxford, UK, third edition (2000).

[107] Naraschewski, M. and Glauber, R. J. "Spatial coherence and density correlations of trapped Bose gases." *Physical Review A* (June 1999). **59**(6), pp. 4595–4607.

[108] Mermin, N. D. and Wagner, H. "Absence of Ferromagnetism or Antiferromagnetism in One- or Two-Dimensional Isotropic Heisenberg Models." *Physical Review Letters* (November 1966). **17**(22), pp. 1133–1136.

[109] Hohenberg, P. C. "Existence of Long-Range Order in One and Two Dimensions." *Physical Review* (June 1967). **158**(2), pp. 383–386.

[110] Hanbury Brown, R. and Twiss, R. Q. "Correlation between Photons in two Coherent Beams of Light." *Nature* (January 1956). **177**(4497), pp. 27–29.

[111] Schellekens, M., Hoppeler, R., Perrin, A., Gomes, J. V., Boiron, D., Aspect, A., and Westbrook, C. I. "Hanbury Brown Twiss Effect for Ultracold Quantum Gases." *Science* (October 2005). **310**(5748), pp. 648–651.

[112] Tonks, L. "The Complete Equation of State of One, Two and Three-Dimensional Gases of Hard Elastic Spheres." *Physical Review* (November 1936). **50**(10), pp. 955–963.

[113] Dunjko, V., Lorent, V., and Olshanii, M. "Bosons in Cigar-Shaped Traps: Thomas-Fermi Regime, Tonks-Girardeau Regime, and In Between." *Physical Review Letters* (June 2001). **86**(24), pp. 5413–5416.

[114] Haller, E., Gustavsson, M., Mark, M. J., Danzl, J. G., Hart, R., Pupillo, G., and Nägerl, H.-C. "Realization of an Excited, Strongly Correlated Quantum Gas Phase." *Science* (2009). **325**(5945), pp. 1224–1227.

[115] Petrov, D. S., Shlyapnikov, G. V., and Walraven, J. T. M. "Regimes of Quantum Degeneracy in Trapped 1D Gases." *Physical Review Letters* (October 2000). **85**(18), pp. 3745–3749.

[116] Hellweg, D., Cacciapuoti, L., Kottke, M., Schulte, T., Sengstock, K., Ertmer, W., and Arlt, J. J. "Measurement of the Spatial Correlation Function of Phase Fluctuating Bose–Einstein Condensates." *Physical Review Letters* (July 2003). **91**(1), p. 010406.

[117] Richard, S., Gerbier, F., Thywissen, J. H., Hugbart, M., Bouyer, P., and Aspect, A. "Momentum Spectroscopy of 1D Phase Fluctuations in Bose-Einstein Condensates." *Physical Review Letters* (July 2003). **91**(1), p. 010405.

[118] Hugbart, M., Retter, J. A., Gerbier, F., Varón, A. F., Richard, S., Thywissen, J. H., Clément, D., Bouyer, P., and Aspect, A. "Coherence length of an elongated condensate: A study by matter-wave interferometry." *The European Physical Journal D* (August 2005). **35**(1), pp. 155–163.

[119] Gallucci, D., Cockburn, S. P., and Proukakis, N. P. "Phase coherence in quasi-condensate experiments: An *ab initio* analysis via the stochastic Gross-Pitaevskii equation." *Physical Review A* (July 2012). **86**(1), p. 013627.

[120] Armijo, J. "Direct Observation of Quantum Phonon Fluctuations in a One-Dimensional Bose Gas." *Physical Review Letters* (June 2012). **108**(22), p. 225306.

[121] Witkowska, E., Deuar, P., Gajda, M., and Rzążewski, K. "Solitons as the Early Stage of Quasicondensate Formation during Evaporative Cooling." *Physical Review Letters* (March 2011). **106**(13), p. 135301.

[122] Karpiuk, T., Deuar, P., Bienias, P., Witkowska, E., Pawłowski, K., Gajda, M., Rzążewski, K., and Brewczyk, M. "Spontaneous Solitons in the Thermal Equilibrium of a Quasi-1D Bose Gas." *Physical Review Letters* (November 2012). **109**(20), p. 205302.

[123] Fontanesi, L., Wouters, M., and Savona, V. "Superfluid to Bose-Glass Transition in a 1D Weakly Interacting Bose Gas." *Physical Review Letters* (July 2009). **103**(3), p. 030403.

[124] Abdullaev, F., Galimzyanov, R., and Ismatullaev, K. "Quasi 1D Bose–Einstein condensate flow past a nonlinear barrier." *Physics Letters A* (October 2012). **376**(45), pp. 3372–3376.

[125] Wick, G. C. "The Evaluation of the Collision Matrix." *Physical Review* (October 1950). **80**(2), pp. 268–272.

[126] Ketterle, W. and van Druten, N. J. "Bose–Einstein condensation of a finite number of particles trapped in one or three dimensions." *Physical Review A* (July 1996). **54**(1), pp. 656–660.

[127] Inouye, S., Andrews, M. R., Stenger, J., Miesner, H.-J., Stamper-Kurn, D. M., and Ketterle, W. "Observation of Feshbach resonances in a Bose–Einstein condensate." *Nature* (March 1998). **392**(6672), pp. 151–154.

[128] Dalibard, J. "Collisional dynamics of ultra-cold atomic gases." In *Proceedings of the International School of Physics "Enrico Fermi," Course CXL*, edited by M. Inguscio, S. Stringari, and C. Wieman. IOS Press, Varenna, Lake Como (1999) pp. 321–349.

[129] Köhler, T., Góral, K., and Julienne, P. S. "Production of cold molecules via magnetically tunable Feshbach resonances." *Reviews of Modern Physics* (December 2006). **78**(4), pp. 1311–1361.

[130] Bolda, E. L., Tiesinga, E., and Julienne, P. S. "Pseudopotential model of ultracold atomic collisions in quasi-one- and two-dimensional traps." *Physical Review A* (September 2003). **68**(3), p. 032702.

[131] Guéry-Odelin, D., Söding, J., Desbiolles, P., and Dalibard, J. "Is Bose-Einstein condensation of atomic cesium possible?" *Europhysics Letters (EPL)* (October 1998). **44**(1), pp. 25–30.

[132] Leo, P. J., Williams, C. J., and Julienne, P. S. "Collision Properties of Ultracold C 133 s Atoms." *Physical Review Letters* (September 2000). **85**(13), pp. 2721–2724.

[133] Davis, K. B., Mewes, M. O., Andrews, M. R., van Druten, N. J., Durfee, D. S., Kurn, D. M., and Ketterle, W. "Bose-Einstein Condensation in a Gas of Sodium Atoms." *Physical Review Letters* (November 1995). **75**(22), pp. 3969–3973.

[134] Ali, D. B., Badr, T., Brézillon, T., Dubessy, R., Perrin, H., and Perrin, A. "Detailed study of a transverse field Zeeman slower." *Journal of Physics B: Atomic, Molecular and Optical Physics* (March 2017). **50**(5), p. 055008.

[135] Steck, D. A. "Sodium D Line Data." available online at `http://steck.us/alkalidata` (revision 2.2.1, 21 November 2019).

[136] Greiner, M., Bloch, I., Mandel, O., Hänsch, T. W., and Esslinger, T. "Exploring Phase Coherence in a 2D Lattice of Bose-Einstein Condensates." *Physical Review Letters* (October 2001). **87**(16), p. 160405.

[137] Bouchoule, I., Morinaga, M., Salomon, C., and Petrov, D. S. "Cesium gas strongly confined in one dimension: Sideband cooling and collisional properties." *Physical Review A* (February 2002). **65**(3), p. 033402.

[138] Moritz, H., Stöferle, T., Köhl, M., and Esslinger, T. "Exciting Collective Oscillations in a Trapped 1D Gas." *Physical Review Letters* (December 2003). **91**(25), p. 250402.

[139] Catani, J., Lamporesi, G., Naik, D., Gring, M., Inguscio, M., Minardi, F., Kantian, A., and Giamarchi, T. "Quantum dynamics of impurities in a one-dimensional Bose gas." *Physical Review A* (February 2012). **85**(2), p. 023623.

[140] Roux, C., Emmert, A., Lupascu, A., Nirrengarten, T., Nogues, G., Brune, M., Raimond, J.-M., and Haroche, S. "Bose-Einstein condensation on a superconducting atom chip." *EPL (Europhysics Letters)* (March 2008). **81**(5), p. 56004.

[141] van Es, J. J. P., Wicke, P., van Amerongen, A. H., Rétif, C., Whitlock, S., and van Druten, N. J. "Box traps on an atom chip for one-dimensional quantum gases." *Journal of Physics B: Atomic, Molecular and Optical Physics* (August 2010). **43**(15), p. 155002.

[142] Boyd, M., Streed, E. W., Medley, P., Campbell, G. K., Mun, J., Ketterle, W., and Pritchard, D. E. "Atom trapping with a thin magnetic film." *Physical Review A* (October 2007). **76**(4), p. 043624.

[143] Alloschery, O., Mathevet, R., Weiner, J., and Lezec, H. J. "All-optical atom surface traps implemented with one-dimensional planar diffractive microstructures." *Optics Express* (2006). **14**(26), p. 12568.

[144] Palzer, S., Zipkes, C., Sias, C., and Köhl, M. "Quantum Transport through a Tonks-Girardeau Gas." *Physical Review Letters* (October 2009). **103**(15), p. 150601.

[145] Estève, J., Schumm, T., Trebbia, J.-B., Bouchoule, I., Aspect, A., and Westbrook, C. I. "Realizing a stable magnetic double-well potential on an atom chip." *The European Physical Journal D* (August 2005). **35**(1), pp. 141–146.

[146] Schumm, T., Hofferberth, S., Andersson, L. M., Wildermuth, S., Groth, S., Bar-Joseph, I., Schmiedmayer, J., and Krüger, P. "Matter-wave interferometry in a double well on an atom chip." *Nature Physics* (October 2005). **1**(1), pp. 57–62.

[147] Reichel, J., Hänsel, W., and Hänsch, T. W. "Atomic Micromanipulation with Magnetic Surface Traps." *Physical Review Letters* (October 1999). **83**(17), pp. 3398–3401.

[148] Dieckmann, K., Spreeuw, R. J. C., Weidemüller, M., and Walraven, J. T. M. "Two-dimensional magneto-optical trap as a source of slow atoms." *Physical Review A* (November 1998). **58**(5), pp. 3891–3895.

[149] Farkas, D. M., Hudek, K. M., Salim, E. A., Segal, S. R., Squires, M. B., and Anderson, D. Z. "A compact, transportable, microchip-based system for high repetition rate production of Bose–Einstein condensates." *Applied Physics Letters* (March 2010). **96**(9), p. 093102.

[150] Monroe, C., Swann, W., Robinson, H., and Wieman, C. "Very cold trapped atoms in a vapor cell." *Physical Review Letters* (September 1990). **65**(13), pp. 1571–1574.

[151] Gibble, K. E., Kasapi, S., and Chu, S. "Improved magneto-optic trapping in a vapor cell." *Optics Letters* (April 1992). **17**(7), p. 526.

[152] Martin, J.-M., Bade, S., Dubosclard, W., Khan, M. A., Kim, S., Garraway, B. M., and Alzar, C. L. G. "Pumping Dynamics of Cold-Atom Experiments in a Single Vacuum Chamber." *Physical Review Applied* (July 2019). **12**(1), p. 014033.

[153] Rodebush, W. H. and Walters, E. G. "The Vapor Pressure and Density of Sodium." *Journal of the American Chemical Society* (July 1930). **52**(7), pp. 2654–2665.

[154] Fiock, E. F. and Rodebush, W. H. "The Vapor Pessures and Thermal Properties of Potassium and some Alkali Halides." *Journal of the American Chemical Society* (October 1926). **48**(10), pp. 2522–2528.

[155] Bohdansky, J. and Schins, H. E. J. "Vapor pressure of different metals in the pressure range of 50 to 4000 torr." *The Journal of Physical Chemistry* (January 1967). **71**(2), pp. 215–217.

[156] Stull, D. R. "Vapor Pressure of Pure Substances. Organic and Inorganic Compounds." *Industrial & Engineering Chemistry* (April 1947). **39**(4), pp. 517–540.

[157] Hobbs, P. C. D. *Building Electro-Optical Systems: Making it all Work*. Wiley Series in Pure and Applied Optics, John Wiley & Sons, Inc., Hoboken, N.J. (July 2009).

[158] Boyd, R. W. *Nonlinear Optics*. Academic Press, San Diego, Calif., fourth edition (March 2020).

[159] Donley, E. A., Heavner, T. P., Levi, F., Tataw, M. O., and Jefferts, S. R. "Double-pass acousto-optic modulator system." *Review of Scientific Instruments* (June 2005). **76**(6), p. 063112.

[160] Lu, B. and Wang, D. "Note: A four-pass acousto-optic modulator system for laser cooling of sodium atoms." *Review of Scientific Instruments* (July 2017). **88**(7), p. 076105.

[161] Saleh, B. E. A. and Teich, M. C. *Fundamentals of Photonics*. Wiley Series in Pure and Applied Optics, Wiley, New York, N.Y. (1991).

[162] Preston, D. W. "Doppler-free saturated absorption: Laser spectroscopy." *American Journal of Physics* (November 1996). **64**(11), pp. 1432–1436.

[163] Juncar, P., Pinard, J., Hamon, J., and Chartier, A. "Absolute Determination of the Wavelengths of the Sodium D_1 and D_2 Lines by Using a CW Tunable Dye Laser Stabilized on Iodine." *Metrologia* (July 1981). **17**(3), pp. 77–79.

[164] Chen, L. *High-Precision Spectroscopy of Molecular Iodine: From Optical Frequency Standards to Global Descriptions of HyperÖne Interactions and Associated Electronic Structure*. PhD Thesis, University of Colorado, Colo. (2005).

[165] Foot, C. J. *Atomic Physics*. Number 7 in Oxford Master Series in Physics, Oxford University Press, Oxford, UK (2005).

[166] Earnshaw, S. "On the nature of the molecular forces which regulate the constitution of the luminiferous ether." In *Cambridge Philosophical Society*, volume 7 (1842) pp. 97–112.

[167] Ashkin, A. and Gordon, J. P. "Stability of radiation-pressure particle traps: an optical Earnshaw theorem." *Optics Letters* (October 1983). **8**(10), p. 511.

[168] Wing, W. H. "On neutral particle trapping in quasistatic electromagnetic fields." *Progress in Quantum Electronics* (January 1984). **8**(3-4), pp. 181–199.

[169] Kartoshkin, V. A. "Spin exchange during collisions of two sodium atoms." *Optics and Spectroscopy* (April 2014). **116**(4), pp. 548–551.

[170] Bergeman, T., Erez, G., and Metcalf, H. J. "Magnetostatic trapping fields for neutral atoms." *Physical Review A* (February 1987). **35**(4), pp. 1535–1546.

[171] Majorana, E. "Atomi orientati in campo magnetico variabile." *Nuovo Cimento* (1932). **9**, pp. 43–50.

[172] Larmor, J. "On the theory of the magnetic influence on spectra; and on the radiation from moving ions." *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* (December 1897). **44**(271), pp. 503–512.

[173] Brink, D. M. and Sukumar, C. V. "Majorana spin-flip transitions in a magnetic trap." *Physical Review A* (September 2006). **74**(3), p. 035401.

[174] Bücker, R., Perrin, A., Manz, S., Betz, T., Koller, C., Plisson, T., Rottmann, J., Schumm, T., and Schmiedmayer, J. "Single-particle-sensitive imaging of freely propagating ultracold atoms." *New Journal of Physics* (October 2009). **11**(10), p. 103039.

[175] Janesick, J. R. *Scientific Charge-coupled Devices*. SPIE Press, Bellingham, Wash. (2001).

[176] Jiang, Y., Jedrkiewicz, O., Minardi, S., DiTrapani, P., Mosset, A., Lantz, E., and Devaux, F. "Retrieval of spatial shot-noise in the full dynamic range of calibrated CCD cameras." *The European Physical Journal D - Atomic, Molecular and Optical Physics* (March 2003). **22**(3), pp. 521–526.

[177] Ham, M., Dae, I., and Choi, C. "LPD: Low Power Display Mechanism for Mobile and Wearable Devices." In *2015 USENIX Annual Technical Conference (USENIX ATC 15)*USENIX Association, Santa Clara, CA (July 2015) pp. 587–598.

[178] Bücker, R., Perrin, A., Manz, S., Betz, T., Koller, C., Plisson, T., Rottmann, J., Schumm, T., and Schmiedmayer, J. "Single-particle-sensitive imaging of freely propagating ultracold atoms." *New Journal of Physics* (2009). **11**(10), p. 103039.

[179] Ashley, C. W. *The Ashley Book of Knots*. Mariners' Museum Publications, Doubleday, New York, N.Y. (1944).

[180] Lefrançois, M., Montuclard, J., and Rouaud, C. "Low-load metal seals to replace elastomer O-rings: the Helicoflex-delta (δ) seals." *Vacuum* (January 1990). **41**(7-9), pp. 1879–1881.

[181] Gonczy, J., Ferry, R., Niemann, R., and Roop, B. "UHV seal studies for the Advanced Photon Source storage ring vacuum system." In *Conference Record of the 1991 IEEE Particle Accelerator Conference*IEEE, San Francisco, CA, USA (1991) pp. 2304–2306.

[182] von Sehren, C., Grulke, O., von Stechow, A., Böttger, L., Edlund, E., and Laube, R. "Ultra high vacuum ZnSe window flange design for Phase Contrast Imaging diagnostics for the Wendelstein 7-X stellarator." In *Proceeding of 30th Symposium on Fusion Technology (SOFT)*EUROfusion, Abingdon, UK (2018) pp. 1–4.

[183] White, G. K. *Experimental Techniques in Low-Temperature Physics*. Number 59 in Monographs on the Physics and Chemistry of Materials, Clarendon Press, Oxford, UK, third edition (1979).

[184] Manz, S., Bücker, R., Betz, T., Koller, C., Hofferberth, S., Mazets, I. E., Imambekov, A., Demler, E., Perrin, A., Schmiedmayer, J., and Schumm, T. "Two-point density correlations of quasicondensates in free expansion." *Physical Review A* (March 2010). **81**(3), p. 031610.

[185] Imambekov, A., Mazets, I. E., Petrov, D. S., Gritsev, V., Manz, S., Hofferberth, S., Schumm, T., Demler, E., and Schmiedmayer, J. "Density ripples in expanding low-dimensional gases as a probe of correlations." *Physical Review A* (September 2009). **80**(3), p. 033604.

[186] Reinaudi, G., Lahaye, T., Wang, Z., and Guéry-Odelin, D. "Strong saturation absorption imaging of dense clouds of ultracold atoms." *Optics Letters* (November 2007). **32**(21), pp. 3143–3145.

[187] Tan, S. "Large momentum part of a strongly correlated Fermi gas." *Annals of Physics* (December 2008). **323**(12), pp. 2971–2986.

[188] Tan, S. "Energetics of a strongly correlated Fermi gas." *Annals of Physics* (December 2008). **323**(12), pp. 2952–2970.

[189] Tan, S. "Generalized virial theorem and pressure relation for a strongly correlated Fermi gas." *Annals of Physics* (December 2008). **323**(12), pp. 2987–2990.

[190] Lang, G., Vignolo, P., and Minguzzi, A. "Tan's contact of a harmonically trapped one-dimensional Bose gas: Strong-coupling expansion and conjectural approach at arbitrary interactions." *The European Physical Journal Special Topics* (May 2017). **226**(7), pp. 1583–1591.

[191] Yao, H., Clément, D., Minguzzi, A., Vignolo, P., and Sanchez-Palencia, L. "Tan's Contact for Trapped Lieb-Liniger Bosons at Finite Temperature." *Physical Review Letters* (November 2018). **121**(22), p. 220402.

[192] Chin, C., Kerman, A. J., Vuletić, V., and Chu, S. "Sensitive Detection of Cold Cesium Molecules Formed on Feshbach Resonances." *Physical Review Letters* (January 2003). **90**(3), p. 033201.

[193] Fedichev, P. O., Reynolds, M. W., and Shlyapnikov, G. V. "Three-Body Recombination of Ultracold Atoms to a Weakly BoundsLevel." *Physical Review Letters* (September 1996). **77**(14), pp. 2921–2924.

[194] Mehta, N. P., Rittenhouse, S. T., D'Incao, J. P., von Stecher, J., and Greene, C. H. "A general theoretical description of N-body recombination." *Physical Review Letters* (October 2009). **103**(15), p. 153201. ArXiv: 0903.4145.

[195] Marte, A., Volz, T., Schuster, J., Dürr, S., Rempe, G., van Kempen, E. G. M., and Verhaar, B. J. "Feshbach Resonances in Rubidium 87: Precision Measurement and Analysis." *Physical Review Letters* (December 2002). **89**(28), p. 283202.

[196] Chin, C., Vuletić, V., Kerman, A. J., Chu, S., Tiesinga, E., Leo, P. J., and Williams, C. J. "Precision Feshbach spectroscopy of ultracold Cs_2." *Physical Review A* (September 2004). **70**(3), p. 032701.

[197] Bragg, W. L. "The diffraction of short electromagnetic waves by a crystal." *Proceedings of the Cambridge Philosophcal Society* (1912). **17**(43).

[198] Martin, P., Oldaker, B., Miklich, A., and Pritchard, D. "Bragg scattering of atoms from a standing light wave." *Physical Review Letters* (February 1988). **60**(6), pp. 515–518.

[199] Stenger, J., Inouye, S., Chikkatur, A. P., Stamper-Kurn, D. M., Pritchard, D. E., and Ketterle, W. "Bragg Spectroscopy of a Bose-Einstein Condensate." *Physical Review Letters* (June 1999). **82**(23), pp. 4569–4573.

[200] Stamper-Kurn, D. M., Chikkatur, A. P., Görlitz, A., Inouye, S., Gupta, S., Pritchard, D. E., and Ketterle, W. "Excitation of Phonons in a Bose-Einstein Condensate by Light Scattering." *Physical Review Letters* (October 1999). **83**(15), pp. 2876–2879.

[201] Kozuma, M., Deng, L., Hagley, E. W., Wen, J., Lutwak, R., Helmerson, K., Rolston, S. L., and Phillips, W. D. "Coherent Splitting of Bose-Einstein Condensed Atoms with Optically Induced Bragg Diffraction." *Physical Review Letters* (February 1999). **82**(5), pp. 871–875.

[202] Zambelli, F., Pitaevskii, L., Stamper-Kurn, D. M., and Stringari, S. "Dynamic structure factor and momentum distribution of a trapped Bose gas." *Physical Review A* (May 2000). **61**(6), p. 063608.

[203] Fabbri, N., Clément, D., Fallani, L., Fort, C., Modugno, M., van der Stam, K. M. R., and Inguscio, M. "Excitations of Bose-Einstein condensates in a one-dimensional periodic potential." *Physical Review A* (April 2009). **79**(4), p. 043623.

[204] Meinert, F., Panfil, M., Mark, M., Lauber, K., Caux, J.-S., and Nägerl, H.-C. "Probing the Excitations of a Lieb-Liniger Gas from Weak to Strong Coupling." *Physical Review Letters* (August 2015). **115**(8), p. 085301.

[205] Caux, J.-S. and Calabrese, P. "Dynamical density-density correlations in the one-dimensional Bose gas." *Physical Review A* (September 2006). **74**(3), p. 031605.

[206] Meyrath, T. P., Schreck, F., Hanssen, J. L., Chuu, C.-S., and Raizen, M. G. "Bose-Einstein condensate in a box." *Physical Review A* (April 2005). **71**(4), p. 041604.

[207] Estève, J., Aussibal, C., Schumm, T., Figl, C., Mailly, D., Bouchoule, I., Westbrook, C. I., and Aspect, A. "Role of wire imperfections in micromagnetic traps for atoms." *Physical Review A* (October 2004). **70**(4), p. 043629.

[208] Folman, R., Krüger, P., Schmiedmayer, J., Denschlag, J., and Henkel, C. "Microscopic Atom Optics: From Wires to an Atom Chip." In *Advances In Atomic, Molecular, and Optical Physics*, volume 48pp. 263–356. Elsevier (2002).

[209] Gott, Y. V., Ioffe, M. S., and Tel'Kovskii, V. G. "Some new results on plasma confinement in a magnetic trap." *reprinted in Nucl. Fusion, Supplement* (1962). **3**, pp. 1045–1047.

[210] Pritchard, D. E. "Cooling Neutral Atoms in a Magnetic Trap for Precision Spectroscopy." *Physical Review Letters* (October 1983). **51**(15), pp. 1336–1339.

[211] Simon, N. J., Drexler, E. S., and Reed, R. P. *Properties of Copper and Copper Alloys at Cryogenic Temperature*. Number 177 in NIST Monographs, U.S. Government Printing Office, Washington DC (February 1992).

[212] Jameson, E. C. *Electrical Discharge Machining*. Society of Manufacturing Engineers, Machining Technology Association, Dearborn, Mich. (2001).

[213] Luiten, O. J., Reynolds, M. W., and Walraven, J. T. M. "Kinetic theory of the evaporative cooling of a trapped gas." *Physical Review A* (January 1996). **53**(1), pp. 381–389.

[214] Ketterle, W. and Druten, N. J. V. "Evaporative Cooling of Trapped Atoms." In *Advances In Atomic, Molecular, and Optical Physics*, volume 37pp. 181–236. Elsevier (1996).

[215] Driscoll, T. A. and Trefethen, L. N. *Schwarz–Christoffel mapping*. Number 8 in Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press, Cambridge, UK (2002).

[216] Horowitz, P. *The Art of Electronics*. Cambridge University Press, New York, N.Y., third edition edition (2015).

[217] Vegt, A. v. d. and Govaert, L. *Polymeren: van keten tot kunststof*. DUP Blue Print, Delft (2003).

[218] Bullis, W. M. "Silicon Material Properties." In *Handbook of Semiconductor Silicon Technology*chapter 6, pp. 347–450. Noyes Publications, Park Ridge, N.J. (1990).

[219] van Es, J. J. P., Whitlock, S., Fernholz, T., van Amerongen, A. H., and van Druten, N. J. "Longitudinal character of atom-chip-based rf-dressed potentials." *Physical Review A* (June 2008). **77**(6), p. 063623.

[220] Perrin, H. and Garraway, B. M. "Trapping Atoms With Radio Frequency Adiabatic Potentials." In *Advances In Atomic, Molecular, and Optical Physics*, volume 66, edited by E. Arimondo, C. C. Lin, and S. F. Yelin, chapter 4, pp. 181–262. Academic Press (2017).

[221] Schumm, T., Estève, J., Figl, C., Trebbia, J.-B., Aussibal, C., Nguyen, H., Mailly, D., Bouchoule, I., Westbrook, C., and Aspect, A. "Atom chips in the real world: the effects of wire corrugation." *The European Physical Journal D* (February 2005). **32**(2), pp. 171–180.

[222] Dellinger, J. "The temperature coefficient of resistance of copper." *Journal of the Franklin Institute* (September 1910). **170**(3), pp. 213–216.

[223] Heisler, S. I. *The Wiley Engineer's Desk Reference: A Concise Guide for the Professional Engineer*. Wiley, New York, N.Y., second edition (1998).

[224] Trebbia, J.-B., Garrido Alzar, C. L., Cornelussen, R., Westbrook, C. I., and Bouchoule, I. "Roughness Suppression via Rapid Current Modulation on an Atom Chip." *Physical Review Letters* (June 2007). **98**(26), p. 263201.

[225] Bücker, R. *Fluorescence Imaging of Ultracold Atoms.* Master's thesis, University of Heidelberg, Heidelberg (2007).

[226] Born, M. and Wolf, E. *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light.* Cambridge University Press, Cambridge, UK, 7th edition (1999).

[227] Shannon, C. E. "Communication in the Presence of Noise." *Proceedings of the IRE* (January 1949). **37**(1), pp. 10–21.

[228] Gomes, J. V., Perrin, A., Schellekens, M., Boiron, D., Westbrook, C. I., and Belsley, M. "Theory for a Hanbury Brown Twiss experiment with a ballistically expanding cloud of cold atoms." *Physical Review A* (November 2006). **74**(5), p. 053607.

[229] Alt, W. "An objective lens for efficient fluorescence detection of single atoms." *Optik* (2002). **113**(3), pp. 142–144.

[230] Gustavson, T. L., Chikkatur, A. P., Leanhardt, A. E., Görlitz, A., Gupta, S., Pritchard, D. E., and Ketterle, W. "Transport of Bose-Einstein condensates with optical tweezers." *Physical Review Letters* (December 2001). **88**(2), p. 020401.

[231] Schrader, D., Kuhr, S., Alt, W., Müller, M., Gomer, V., and Meschede, D. "An optical conveyor belt for single neutral atoms." *Applied Physics B* (2001). **73**, pp. 819–824.

[232] Couvert, A., Kawalec, T., Reinaudi, G., and Guéry-Odelin, D. "Optimal transport of ultracold atoms in the non-adiabatic regime." *EPL (Europhysics Letters)* (2008). **83**(1), p. 13001.

[233] Lewandowski, H., Harber, D., Whitaker, D., and Cornell, E. "Simplified system for creating a Bose-Einstein condensate." *Journal of Low Temperature Physics* (2003). **132**(5/6), p. 309.

[234] Händel, S., Marchant, A. L., Wiles, T. P., Hopkins, S. A., and Cornish, S. L. "Magnetic transport apparatus for the production of ultracold atomic gases in the vicinity of a dielectric surface." *Review of Scientific Instruments* (January 2012). **83**(1), p. 013105.

[235] Schmid, S., Thalhammer, G., Winkler, K., Lang, F., and Denschlag, J. H. "Long distance transport of ultracold atoms using a 1D optical lattice." *New Journal of Physics* (August 2006). **8**(8), pp. 159–159.

[236] Hänsel, W., Hommelhoff, P., Hänsch, T., and Reichel, J. "Bose-Einstein condensation on a microelectronic chip." *Nature* (2001). **413**, pp. 498–501.

[237] Hänsel, W., Reichel, J., Hommelhoff, P., and Hänsch, T. "Magnetic conveyor belt for transporting and merging trapped atom clouds." *Physical Review Letters* (2001). **86**(4), pp. 608–611.

[238] Lahaye, T., Reinaudi, G., Wang, Z., Couvert, A., and Guéry-Odelin, D. "Transport of atom packets in a train of Ioffe-Pritchard traps." *Physical Review A* (September 2006). **74**(3), p. 033622.

[239] Minniberger, S., Diorico, F., Haslinger, S., Hufnagel, C., Novotny, C., Lippok, N., Majer, J., Koller, C., Schneider, S., and Schmiedmayer, J. "Magnetic conveyor belt transport of ultracold atoms to a superconducting atomchip." *Applied Physics B* (September 2014). **116**(4), pp. 1017–1021.

[240] Rio Fernandes, D. *Trapping and cooling of fermionic alkali atoms to quantum degeneracy. Sub-Doppler cooling of Potassium-40 and Lithium-6 in gray molasses*. Ph.D. thesis, Université Pierre et Marie Curie, Paris (September 2014).

[241] Schaff, J.-F., Song, X.-L., Capuzzi, P., Vignolo, P., and Labeyrie, G. "Shortcut to adiabaticity for an interacting Bose–Einstein condensate." *EPL (Europhysics Letters)* (January 2011). **93**(2), p. 23001.

[242] Torrontegui, E., Ibáñez, S., Chen, X., Ruschhaupt, A., Guéry-Odelin, D., and Muga, J. G. "Fast atomic transport without vibrational heating." *Physical Review A* (January 2011). **83**(1), p. 013415.

[243] Schaff, J.-F., Capuzzi, P., Labeyrie, G., and Vignolo, P. "Shortcuts to adiabaticity for trapped ultracold gases." *New Journal of Physics* (November 2011). **13**(11), p. 113017.

[244] Zhang, Q., Chen, X., and Guéry-Odelin, D. "Fast and optimal transport of atoms with nonharmonic traps." *Physical Review A* (October 2015). **92**(4), p. 043410.

[245] Zhang, Q., Muga, J. G., Guéry-Odelin, D., and Chen, X. "Optimal shortcuts for atomic transport in anharmonic traps." *Journal of Physics B: Atomic, Molecular and Optical Physics* (June 2016). **49**(12), p. 125503.

[246] Badr, T., Ali, D. B., Seaward, J., Guo, Y., Wiotte, F., Dubessy, R., Perrin, H., and Perrin, A. "Comparison of time profiles for the magnetic transport of cold atoms." *Applied Physics B* (2019). **125**(6), p. 102.

[247] Jovcic, D. "Bidirectional, high-power DC transformer." *IEEE Transactions on Power Delivery* (October 2009). **24**(4), pp. 2276–2283.

[248] Zhang, J. *Bidirectional DC-DC power converter design optimization, modeling and control*. Ph.D. thesis, Virginia Tech, Va. (2008).

[249] Tiesinga, E., Williams, C., Julienne, P., Jones, K., Lett, P., and Phillips, W. "A spectroscopic determination of scattering lengths for sodium atom collisions." *Journal of Research of the National Institute of Standards and Technology* (July 1996). **101**(4), p. 505.

[250] Seaward, J. *Extending Bethe Ansatz Solutions to the Non-Crystallographic Reflection Groups H_3 and H_4*. Master's thesis, UMass, Boston, Boston, Mass. (August 2017).

[251] Caux, J.-S. and Mossel, J. "Remarks on the notion of quantum integrability." *Journal of Statistical Mechanics: Theory and Experiment* (February 2011). **2011**(02), p. P02023.

[252] Calogero, F. "Exactly solvable one-dimensional many-body problems." *Lettere al Nuovo Cimento (1971-1985)* (July 1975). **13**(11), pp. 411–416.

[253] Eisert, J., Friesdorf, M., and Gogolin, C. "Quantum many-body systems out of equilibrium." *Nature Physics* (February 2015). **11**(2), pp. 124–130.

[254] Rigol, M., Dunjko, V., and Olshanii, M. "Thermalization and its mechanism for generic isolated quantum systems." *Nature* (April 2008). **452**(7189), pp. 854–858.

[255] Redner, S. "A billiard-theoretic approach to elementary one-dimensional elastic collisions." *American Journal of Physics* (December 2004). **72**(12), pp. 1492–1498.

[256] Hwang, Z., Cao, F., and Olshanii, M. "Traces of integrability in relaxation of one-dimensional two-mass mixtures." *J. Stat. Phys.* (2015). **161**, p. 467.

[257] Wang, J., Casati, G., and Prosen, T. "Non-ergodicity and localization of invariant measure for two colliding masses." (2013). ArXiv: 1309.7617.

[258] Gutkin, E. "Billiard dynamics: An updated survey with the emphasis on open problems." *Chaos: An Interdisciplinary Journal of Nonlinear Science* (2012). **22**(2), p. 026116.

[259] Huveneers, F. "Subdiffusive behavior generated by irrational rotations." *Ergodic Theory and Dynamical Systems* (2009). **29**(4), pp. 1217–1233.

[260] Chernov, N. and Markarian, R. *Chaotic Billiards*, volume 127 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, R.I. (July 2006).

[261] Einsiedler, M. and Ward, T. *Ergodic Theory: with a view towards Number Theory*. Springer London, London (2011).

[262] Titulaer, E. "Ergodic features of harmonic-oscillator systems. I." *Physica* (December 1973). **70**(2), pp. 257–275.

[263] Teschl, G. *Ordinary Differential Equations and Dynamical Systems*. Number 140 in Graduate Studies in Mathematics, American Mathematical Society, Providence, R.I. (2012).

[264] Landau, L. D. and Lifshitz, E. M. *Mechanics*, volume 1 of *Course of Theoretical Physics*. Pergamon Press, Oxford, UK, second edition (1969).

[265] de la Llave, R. "A tutorial on KAM theory." In *Proceedings of Symposia in Pure Mathematics*, volume 69, edited by A. Katok, R. de la Llave, Y. Pesin, and H. Weiss, pp. 175–292. American Mathematical Society, Providence, R.I. (2001).

[266] Gutkin, E. "Billiards in polygons." *Physica D: Nonlinear Phenomena* (April 1986). **19**(3), pp. 311–333.

[267] Dragovic, V. and Radnovic, M. "Bifurcations of Liouville tori in elliptical billiards." *Regular and Chaotic Dynamics* (August 2009). **14**(4-5), pp. 479–494. ArXiv: 0902.4233.

[268] Sinaï, Y. G. "Dynamical systems with elastic reflections. Ergodic properties of dispersing billiards." *Uspekhi Matematicheskikh Nauk* (1970). **25**(2), pp. 141–192.

[269] Robnik, M. and Berry, M. V. "Classical billiards in magnetic fields." *Journal of Physics A: Mathematical and General* (June 1985). **18**(9), pp. 1361–1378.

[270] Tabachnikov, S. and Dogru, F. "Dual Billiards." *The Mathematical Intelligencer* (September 2005). **27**(4), pp. 18–25.

[271] Moser, J. "Is the solar system stable?" *The Mathematical Intelligencer* (June 1978). **1**(2), pp. 65–71.

[272] Gutkin, E. "Billiard dynamics: An updated survey with the emphasis on open problems." *Chaos: An Interdisciplinary Journal of Nonlinear Science* (June 2012). **22**(2), p. 026116.

[273] Smillie, J. "The Dynamics of Billiard Flows in Rational Polygons." In *Dynamical Systems, Ergodic Theory and Applications*, volume 100 of *Encyclopaedia of Mathematical Sciences*pp. 360–382. Springer Science & Business Media, Berlin (2000).

[274] Vorobets, Y. B. "Ergodicity of billiards in polygons." *Sbornik: Mathematics* (April 1997). **188**(3), pp. 389–434.

[275] Cipra, B., Hanson, R. M., and Kolan, A. "Periodic trajectories in right-triangle billiards." *Physical Review E* (Aug 1995). **52**, pp. 2066–2071.

[276] Galperin, G. and Zvonkine, D. "Periodic billiard trajectories in right triangles and right-angled tetrahedra." *Regular and Chaotic Dynamics* (2003). **8**(1), pp. 29–44.

[277] Hooper, W. P. "Periodic billiard paths in right triangles are unstable." *Geometriae Dedicata* (2007). **125**(1), pp. 39–46.

[278] Coxeter, H. S. M. *Regular Polytopes*. Dover, New York, N.Y., third edition (1973).

[279] Jeans, J. H. *The Mathematical Theory of Electricity and Magnetism*. University Press, Cambridge, UK, second edition (1911).

[280] Honerkamp, J. *Stochastic Dynamical Systems: Concepts, Numerical Methods, Data Analysis*. Wiley-VCH, New York, N.Y., first edition (2009).

[281] Cristadoro, G., Knight, G., and Esposti, M. D. "Follow the fugitive: an application of the method of images to open dynamical systems." *Journal of Physics A: Mathematical and Theoretical* (July 2013). **46**(27), p. 272001. ArXiv: 1212.0673.

[282] Jepsen, D. W. and Hirschfelder, J. O. "Set of co-ordinate systems which diagonalize the kinetic energy of relative motion." *Proceedings of the National Academy of Sciences* (1959). **45**(2), pp. 249–256.

[283] Olshanii, M., Jacobs, K., Rigol, M., Dunjko, V., Kennard, H., and Yurovsky, V. A. "An exactly solvable model for the integrability-chaos transition in rough quantum billiards." *Nature Communications* (2012). **3**, p. 641.

[284] Araújo Lima, T., Rodríguez-Pérez, S., and de Aguiar, F. M. "Ergodicity and quantum correlations in irrational triangular billiards." *Physical Review E* (June 2013). **87**(6), p. 062902.

[285] Scoquart, T., Seaward, J. J., Jackson, S. G., and Olshanii, M. "Exactly solvable quantum few-body systems associated with the symmetries of the three-dimensional and four-dimensional icosahedra." *SciPost Physics* (October 2016). **1**(1), p. 005. ArXiv: 1608.04402.

[286] Sěba, P. "Wave chaos in singular quantum billiard." *Physical Review Letters* (April 1990). **64**(16), pp. 1855–1858.

[287] McGuire, J. B. and Dirk, C. "Extending the Bethe Ansatz: The Quantum Three-Particle Ring." *Journal of Statistical Physics* (February 2001). **102**(3-4), pp. 971–980.

[288] Frączek, K. and Lemańczyk, M. "On mild mixing of special flows over irrational rotations under piecewise smooth functions." *Ergodic Theory and Dynamical Systems* (2006). **26**(3), pp. 719–738.

[289] Elliott, G. A. and Evans, D. E. "The Structure of the Irrational Rotation C * -Algebra." *The Annals of Mathematics* (November 1993). **138**(3), p. 477.

[290] Simon, B. "Almost periodic Schrödinger operators: A Review." *Advances in Applied Mathematics* (December 1982). **3**(4), pp. 463–490.

[291] Last, Y. "Almost everything about the almost Mathieu operator. I." In *XIth International Congress of Mathematical Physics (Paris, 1994)* (1994) pp. 366–372.

[292] Coxeter, H. S. M. *Introduction to Geometry*. Wiley Classics Library, Wiley, New York, N.Y., second edition (1989).

[293] Masur, H. and Tabachnikov, S. "Rational billiards and flat structures." In *Handbook of Dynamical Systems*, volume 1, part A, edited by B. Hasselblatt and A. Katok, chapter 13, pp. 1015–1089. North-Holland, Amsterdam (2002).

[294] Hardy, G. H. and Wright, E. M. *An Introduction to the Theory of Numbers*. Oxford University Press, Oxford, UK, 6th edition (2008).

[295] Keane, M. "Irrational rotations and quasi-ergodic measures." *Publications mathématiques et informatique de Rennes* (1970). **1**(1), pp. 17–26.

[296] Lang, S. "Transcendental numbers and diophantine approximations." *Bulletin of the American Mathematical Society* (September 1971). **77**(5), pp. 635–678.

[297] Burton, R. and Denker, M. "On the central limit theorem for dynamical systems." *Transactions of the American Mathematical Society* (1987). **302**(2), pp. 715–726.

[298] Weyl, H. "Über die Gleichverteilung von Zahlen mod. Eins (on the distribution of numbers mod one)." *Mathematische Annalen* (September 1916). **77**(3), pp. 313–352.

[299] Montgomery, H. L. "Harmonic Analysis as found in Analytic Number Theory." In *Twentieth Century Harmonic Analysis — A Celebration*, edited by J. S. Byrnes, pp. 271–293. Springer Netherlands, Dordrecht (2001).

[300] Shallit, J. O. "Simple continued fractions for some irrational numbers, II." *Journal of Number Theory* (April 1982). **14**(2), pp. 228–231.

[301] Niven, I. M. *Diophantine Approximations*. Dover Books on Mathematics, Dover Publications, Mineola, N.Y. (2008).

[302] Cuyt, A., Jones, W. B., Petersen, V. B., Verdonk, B., and Waadeland, H. *Handbook of Continued Fractions for Special Functions*. Springer Science & Business Media, Dordrecht (2008).

[303] Bailey, D., Borwein, P., and Plouffe, S. "On the rapid computation of various polylogarithmic constants." *Mathematics of Computation* (April 1997). **66**(218), pp. 903–914.

[304] Kim, D. H. and Seo, B. K. "The waiting time for irrational rotations." *Nonlinearity* (September 2003). **16**(5), pp. 1861–1868.

[305] Falconer, K. J. *Fractal Geometry: Mathematical Foundations and Applications.* Wiley, Chichester, UK (1990).

[306] Li, J., Du, Q., and Sun, C. "An improved box-counting method for image fractal dimension estimation." *Pattern Recognition* (November 2009). **42**(11), pp. 2460–2469.

[307] Hutchinson, J. E. "Fractals and Self Similarity." *Indiana University Mathematics Journal* (October 1981). **30**(5), pp. 713–747.

[308] Hwang, Z., Marx, C. A., Seaward, J., Jitomirskaya, S., and Olshanii, M. "Growth of the Wang-Casati-Prosen counter in an integrable billiard." (November 2020). ArXiv: 2011.09467, submitted to *Journal of Statistiacal Physics*, Nov. 2020.

[309] Casati, G. and Ford, J. "Computer study of ergodicity and mixing in a two-particle, hard point gas system." *Journal of Computational Physics* (1976). **20**(1), pp. 97–109.

[310] Artuso, R., Casati, G., and Gburuarneri, I. "Numerical study on ergodic properties of triangular billiards." *Physical Review E* (1997). **55**, pp. 6384–6390.

[311] Wang, J., Casati, G., and Prosen, T. "Nonergodicity and localization of invariant measure for two colliding masses." *Physical Review E* (April 2014). **89**, p. 042918.

[312] Horvat, M., Esposti, M. D., Isola, S., Prosen, T., and Bunimovich, L. "On ergodic and mixing properties of the triangle map." *Physica D* (March 2009). **238**(4), pp. 395–415.

[313] Casati, G. and Prosen, T. "Triangle map: A model of quantum chaos." *Physical Review Letters* (November 2000). **85**, pp. 4261–4264.

[314] Böhi, P., Riedel, M. F., Hänsch, T. W., and Treutlein, P. "Imaging of microwave fields using ultracold atoms." *Applied Physics Letters* (August 2010). **97**(5), p. 051101.

Résumé

Cette thèse traite d'une expérience visant à confiner des atomes de sodium dans un piège magnétique très anisotrope créé par une puce à atomes. La puce contient un guide micro-onde : nous prévoyons en effet d'utiliser une résonance de Feshbach dans le domaine micro-onde, encore jamais observée, pour contrôler la force des interactions entre atomes et explorer les régions du diagramme de phase d'un gaz unidimensionnel (1D). Un gaz 1D présente des corrélations dans ses fluctuations de densité. Pour les observer, il faut accumuler des données de nombreuses réalisations du gaz, ainsi l'expérience est conçue pour atteindre des taux de répétition élevés. Pour cela, la zone de préparation initiale du nuage est placée à distance de l'enceinte à vide contenant la puce atomique et le nuage froid est transporté via une chaîne de bobines de piégeage magnétique. Nous avons optimisé le transport dans ces bobines afin de pouvoir transporter un nuage sur 30 cm en 600 ms tout en conservant plus de 80% des atomes. Les données sur le gaz sont obtenues par imagerie. Nous présentons trois objectifs pour imager le nuage dans trois conditions différentes : avant, pendant et après la phase de piégeage par la puce.

Une deuxième partie présente des travaux théoriques sur les "compteurs de parité" pour les systèmes dynamiques, permettant d'établir le caractère ergodique du système. Ce travail s'inspire d'un système de billards triangulaires, qui peuvent coder certains systèmes 1D. Nous étudions le comportement d'un compteur pour un système plus simple, les rotations irrationnelles, et trouvons les propriétés de croissance que l'on peut attendre pour tous les compteurs de parité.

Mots-clefs : Atomes froids, résonance de Feshbach, transport magnétique, puce à atomes, gaz unidimensionnels, systèmes dynamiques, théorie ergodique, rotations irrationnelles

Abstract

This thesis treats an experiment which aims to confine sodium atoms in an elongated magnetic trap generated by an atom chip. The chip contains a microwave guide: we intend to find and use a novel Feshbach resonance to tune the atoms' interaction strength and explore the different regions of a one-dimensional (1-D) gases' phase diagram. A 1-D gas is characterized by correlations in density fluctuations. To see these, we must take data over many samples of the gas, so steps were taken to increase the repetition rate. To that end, the initial cooling of the gas takes place in a chamber far from the atom chip and a cold cloud is transported via a chain of magnetic trap coils. We have optimized the transport in these coils so that we can transport a cloud 30 cm in 600 ms while keeping over 80% of the atoms. Data about the gas is obtained via imaging. This thesis presents three objectives for imaging the cloud in three different conditions: before, during, and after it is in the trap generated by the chip.

The second part presents preliminary theoretical work on "parity counters" for dynamical systems. The goal is to use these counters to answer questions of ergodicity. This work was inspired by triangular billiards which can encode 1-D systems. We study a counter for a simpler system, irrational rotations, and find properties of the growth that can be expected in all parity counters.

Keywords: Cold atoms, Feshbach resonance, magnetic transport, atom chip, one-dimensional gases, dynamical systems, ergodic theory, irrational rotations.