

## Durham E-Theses

---

# *Caught In A Trap: Monte Carlo Methods For A Travelling-Wave Zeeman Decelerator*

PAUL ARBUTHNOT WALKER

### How to cite:

---

WALKER, PAUL ARBUTHNOT (2023) Caught In A Trap: Monte Carlo Methods For A Travelling-Wave Zeeman Decelerator. Doctoral thesis, Durham University.

### Use policy

---

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a <https://etheses.durham.ac.uk/id/eprint/15405/> is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

# **Caught In A Trap: Monte Carlo Methods For A Travelling-Wave Zeeman Decelerator**

**Paul Arbuthnot Walker**

---

A thesis presented for the degree of Doctor of  
Philosophy



Department of Physics

Durham University

April, 2023



# Abstract

Our group has previously designed and constructed a modular travelling-wave Zeeman decelerator [1, 2], incorporating a pulsed-valve supersonic expansion source with dielectric barrier discharge. The decelerator comprises a set of custom-made flattened helical coils driven by a bespoke programmable power electronics control system. The decelerator produces very strong longitudinal magnetic forces that generate an effectively three-dimensional moving trapping potential, compared to a ‘conventional’ Zeeman decelerator. An additional transverse focussing field was originally added by a wire quadrupole circuit, with a possible replacement being constructed from permanent magnets. The decelerator is intended as a loading stage for a novel hybrid magnetic trap / magneto-optical trap, designed for the sympathetic cooling of atomic or molecular species for which conventional laser cooling is not an option.

This thesis describes some proposed adjustments and enhancements to the experimental apparatus. A suite of software tools has been developed in order to simulate the various components of the experiment. By its nature the travelling-wave decelerator produces a complex signal at a detector and Monte Carlo analysis has been necessary to understand it. For each part of our simulation process the theoretical underpinning and physical implementation is detailed. Models for fitting and characterising supersonic beam data are developed and demonstrated, and applied to the generation of simulated expansions. Key to the operation of the dynamical simulations is a novel interpolation library which is used to solve particle motion in time-dependent magnetic fields. Computer codes for calculating Stern-Gerlach type forces in different species are described. The phase-space acceptance of the decelerator and proposed trap are estimated.

Metastable argon has been guided and decelerated for the purposes of characterising the decelerator; data is presented and analysed, showing a demonstrable reduction by  $\approx 6\%$  of the kinetic energy of a portion of the beam. Predictions show that this could be increased to  $\approx 14\%$  for the decelerator in its current state, and up to  $\approx 30\%$  should be achievable for the case of argon with a proposed extension of the deceleration stage.

Further predictions are made regarding the deceleration of other species of scientific interest, including a dual-species beam of hydrogen and lithium. Some preliminary results regarding the suggested hybrid trap design are given.



# Declaration

The work in this thesis is based on research carried out at the Department of Physics, the University of Durham, United Kingdom, between October 2017 and March 2023.

It is all my own work unless referenced to the contrary in the text. I confirm that no part of the material offered has previously been submitted by myself for a degree in this or any other University.

Paul A. Walker  
Crook, April, 2023

The copyright of this thesis rests with the author. No quotation from it should be published without their prior written consent and information derived from it should be acknowledged.



*Pereant qui ante nos nostra dixerunt.*

Aelius Donatus

*The computer can't tell you the emotional story. It can give  
you the exact mathematical design, but what's missing is the  
eyebrows.*

Frank Zappa

*...and wrote another chapter  
From the Book of Bad Decisions I have made...*

Clutch



## Acknowledgements

Some years ago I heard someone say that doing a PhD is like doing a group project where your partner is an idiot, but it's your past self. At the time I laughed, but the 'joke' is wearing pretty thin by now...

I suppose I have to start by once again thanking my supervisor, Dr David Carty, for giving me the opportunity to continue working on this fascinating project. I say 'suppose' because it may have seemed at times to be the biggest mistake I ever made, but I came asking for it so David can't take the blame. Cheers chief.

Like most labs, we owe a huge debt of thanks to the mechanical, electronic and research support technicians. Possibly even more vital to any success I've achieved are Lynn and Catherine who were unfailingly cheerful in supplying tea and coffee, and didn't judge me too harshly when I had no cash and had to go in The Book. It's also not hyperbole to say that this work owes a huge debt to The Victoria and The Station House, the two best pubs within many miles of Durham. Indeed, most of it was written up in one or the other of them.

I must thank my two predecessors, firstly Dr Arin Mizouri, who unwittingly taught me the importance of the 'appeal to authority' logical fallacy. Her thesis has proved to be a bottomless pit of interest. And what can I say about Dr Lewis McArd; electronics guru, robotics wrangler, tractor fettle, pool shark, drinking buddy and good friend. What a guy. I absolutely wouldn't have gotten through this without his assistance. I must also acknowledge the debt I owe to Prof Ifan Hughes, who has influenced my time in AtMol / QLM via his graduate course, his being on my review team, his steadfast willingness to go for a pint and a pickled egg and listen to my questions with only a modicum of well-meaning (?) scorn, and to his kindly agreeing to proofread several parts of this work - albeit as an avoidance tactic for the hundred other more important things he has to do. I can be this effusive, safe in the knowledge that he will not become insufferably big-headed because he will never see this. See you all at the 'Friday seminar'; first round is on me.

Lastly but most importantly, mijn lieve Martina, who bears much of the blame for egging me on to enroll on this foolhardy endeavour. If we hadn't both been postgraduates at the same time I don't know how we'd have survived, but 'misery loves company' as they say. Although I have grumbled and moaned all the way I am grateful for you dragging me out to get some 'fresh' air (as fresh as it gets around the ponies). Ik houd van jou, nu schiet op en maak *jouw* thesis af...



# Contents

<b>1</b>	<b>Introduction and Motivation</b>	<b>1</b>
1.1	Techniques . . . . .	1
1.2	Motivation . . . . .	7
1.3	Aims Of This Project And Thesis . . . . .	9
<b>2</b>	<b>Decelerator Experimental Theory and Design</b>	<b>10</b>
2.1	Experimental Layout . . . . .	10
2.2	Beam Source . . . . .	12
2.2.1	Supersonic Expansion . . . . .	12
2.2.2	Even-Lavie Valve . . . . .	14
2.2.3	Dielectric Barrier Discharge . . . . .	15
2.2.4	Stability of the Source . . . . .	16
2.3	Zeeman Deceleration . . . . .	18
2.3.1	The Zeeman Effect . . . . .	18
2.3.2	Zeeman Decelerators . . . . .	22
2.3.3	Travelling-Wave Zeeman Decelerators . . . . .	24
2.3.4	Decelerator Coils . . . . .	26
2.3.5	Decelerator Electronics . . . . .	29
2.3.6	Pulse-Width Modulation . . . . .	29
2.3.7	Quadrupole . . . . .	32
2.4	Beam Detection . . . . .	37
2.4.1	Time-of-Flight Profiles . . . . .	37
2.4.2	Microchannel Plate Detector . . . . .	39
2.4.3	Buffer Amplifier . . . . .	40
2.5	Experimental Parameters . . . . .	43
2.6	Chapter Summary . . . . .	44
<b>3</b>	<b>Data Analysis and Simulation of the Zeeman Decelerator</b>	<b>45</b>
3.1	Analysing Time-of-Flight Data . . . . .	45
3.1.1	ToF at a Detector . . . . .	45
3.1.2	ToF at Dual Detectors . . . . .	50
3.2	Supersonic Expansion Source . . . . .	51
3.3	Arbitrary Function Sampling . . . . .	58
3.4	Magnetic Trapping . . . . .	60
3.5	Source State Population of Argon . . . . .	64

3.6	Calculating Magnetic Fields . . . . .	65
3.6.1	Decelerator Fields . . . . .	65
3.6.2	Static Fields . . . . .	67
3.7	Interpolation . . . . .	71
3.7.1	Spline Interpolation . . . . .	71
3.7.2	Tricubic Interpolation . . . . .	72
3.7.3	Quadcubic Interpolation . . . . .	73
3.7.4	Software Implementation . . . . .	74
3.7.5	Accuracy Of The Tricubic Method . . . . .	76
3.8	Numerical Integration . . . . .	80
3.8.1	Runge-Kutta . . . . .	80
3.8.2	Velocity Verlet . . . . .	81
3.9	Point Location . . . . .	83
3.10	Monte Carlo Simulation of the Decelerator . . . . .	85
<b>4</b>	<b>Applying the Simulations</b>	<b>99</b>
4.1	Decelerator Loading . . . . .	99
4.1.1	Delay Calculation . . . . .	99
4.1.2	Longitudinal Acceptance Model . . . . .	102
4.1.3	Longitudinal Decelerator Loading Model . . . . .	105
4.2	PWM Effects . . . . .	108
4.3	Static Acceptance Model . . . . .	114
4.4	Dynamic Acceptance Model . . . . .	118
4.5	Extending the Decelerator Simulations . . . . .	123
4.5.1	Argon . . . . .	123
4.5.2	Assessing Quadrupole Effects . . . . .	130
4.6	Comparison to Other Travelling-Wave Zeeman Decelerators . . . . .	137
4.6.1	Paris Decelerator . . . . .	137
4.6.2	Tel Aviv Decelerator . . . . .	142
4.6.3	Basel Decelerator . . . . .	146
4.6.4	Conclusion . . . . .	150
<b>5</b>	<b>Outlook</b>	<b>151</b>
5.1	Extending the Decelerator . . . . .	151
5.2	Magnetic Trap / Magneto-Optical Trap . . . . .	153
5.3	Sympathetic Cooling . . . . .	154
5.3.1	Calcium Monohydride plus Lithium . . . . .	155
5.3.2	Co-Deceleration and Trapping of CaH and Li . . . . .	155
5.3.3	Hydrogen plus Lithium . . . . .	161
5.3.4	Co-Deceleration and Trapping of H and Li . . . . .	161

<i>CONTENTS</i>	xiii
5.4 Buffer Gas Source . . . . .	169
5.5 Non-Adiabatic Losses . . . . .	169
<b>6 Conclusion</b>	<b>172</b>
<b>A Additional Argon Data</b>	<b>174</b>



# Chapter 1

## Introduction and Motivation

Experiments utilising cold and ultracold atoms have found application in such fields as metrology [3, 4], quantum simulation [5], and low-temperature chemistry [6, 4]. They are also of increasing interest in the search for new physics beyond the standard model [7]. Generally speaking, ‘cold’ means temperatures in the range  $\approx 1$  mK to 1 K, and ‘ultracold’  $<1$  mK. The concept of temperature being used in this context is the one given by thermodynamics; the ‘translational temperature’ is a parameter which describes the distribution of particle velocities in a system. Technically speaking, this definition of temperature requires that the system be in equilibrium, *i.e.* the particles are frequently colliding with each other so that energy is redistributed, and this is not the case under many experimental conditions - the work in this thesis is concerned with supersonic expansions, in which we assume that the atoms or molecules are of low enough density that thermal contact essentially ceases. Nonetheless, the temperature remains a useful metric for the spread of kinetic energies of the particles of interest.

Utilising molecules instead of atoms could greatly increase the scope of these areas of research [8, 9], but the trapping and cooling of molecules entails additional challenges. Compared to atoms, molecules possess rotational and vibrational degrees of freedom, with respective typical energies of order  $10^{-2}$  and  $10^{-1}$  eV [10]. One possible route to obtaining samples of low-temperature molecules is the Zeeman decelerator, which uses time-varying inhomogeneous magnetic fields to manipulate beams of paramagnetic particles. Our group has designed and constructed a Zeeman decelerator of the ‘travelling-wave’ or ‘moving-trap’ type [1, 2]. This experiment was originally motivated by efforts to produce quantum simulators from arrays of trapped polar molecules, but could have applications in the previously-mentioned fields of controlled chemistry and low-temperature precision spectroscopy. This thesis presents experimental results taken when characterising the experiment, with a comprehensive set of computer simulations developed to interpret these results.

### 1.1 Techniques

This section will briefly describe some of the key techniques used in the creation of cold or ultracold molecules. Roughly speaking, we can distinguish between two types of method; direct, in which a sample of a molecular species of interest is somehow cooled, or indirect, in which the molecules are created from their component atoms or molecules, which are themselves first cooled.

Buffer gas techniques produce cold atoms or molecules by injecting them into a cell containing a cryogenically-cooled carrier such as helium, with typical temperatures between a few hundred mK to a few K [11]. The species of interest is cooled by thermal contact with the buffer. The buffer gas cell may be used to produce a beam of particles. As well as being internally cold, buffer gas sources typically produce beams that are relatively slow in the laboratory frame [12]; for example, Truppe *et*

*al.* described a buffer source of CaF molecules with a flow velocity of  $\approx 150 \text{ m s}^{-1}$  [13]. Another common technique is the supersonic expansion (see below), for which typical flow velocities are likely to be two or three times larger. Where applicable, entrained species may be guided or trapped by electric [14] or magnetic [15] fields. Such guiding offers further control over the population of the beam via velocity selection. A range of molecular species have been produced in cryogenic buffer gas beams, such as CaF, MgF, YbF and NO [16]. Recent advances in computer modelling of such beams have aided in experimental design [17].

As mentioned above, another common technique for the production of beams of cooled atoms or molecules is the supersonic expansion. When expanding into a vacuum, particles undergo a supersonic expansion if the mean free path between collisions  $\lambda_0$  is much smaller than  $D$ , the diameter of the nozzle or aperture [18]. During the expansion there is a rapid redistribution of kinetic energy, giving rise to a new velocity distribution with a high bulk motion in the laboratory frame but a much reduced spread of longitudinal velocities. In the classical view the variance of the velocity distribution defines the (translational) temperature of the beam. Strictly speaking this is not applicable post-expansion since the density drops and collisions cease, meaning that the beam is not in thermodynamic equilibrium; however, temperature is a useful and commonly-used metric when describing a beam. The source or ‘stagnation’ chamber is typically cryogenically cooled since it can be shown that the terminal flow velocity  $v_T$  and terminal temperature  $T_T$  are both proportional to the stagnation temperature  $T_S$  [19]. More detail is given in section 2.2.1.

Although supersonic expansion is therefore a useful tool for making samples of ‘cold’ particles, the fact that they have fast bulk motion must be accounted for. One approach is to work in the moving frame, in which the particles have low-energy collisions. For example, Amarasinghe and Suits describe studying collisions between ground-state and Rydberg-state Xe atoms in a single beam [20]. Similarly, a beam may be produced that contains multiple species; Perreault *et al.* studied collisions in a collimated supersonic expansion of HD and He [21]. A standard technique in studies of ion collisions is the merging of beams, with one or both guided with electrostatic or magnetic fields [22]. Wei *et al.* proposed utilising this method for beams of cold neutral particles [23]. Henson *et al.* merged beams of Ar and H<sub>2</sub> with metastable He [24].

Alternatively to working in a moving frame, one can attempt to remove the bulk motion of a supersonic expansion. A mechanical method is to use a rotating nozzle to produce the expansion, such that the rotation counters the bulk motion of the beam, first demonstrated by Gupta and Herschbach [25, 26]. A related system is the centrifuge decelerator [27]. A different technique to produce a slow sample from a molecular beam is to decelerate it. Stark and Zeeman decelerators respectively use time-varying electric and magnetic inhomogeneous fields to manipulate beams of atoms or molecules. The deceleration process does not ‘cool’ a molecular beam in the sense of reducing the spread of energies of its degrees of freedom, but rather reduces the bulk motion of a packet of particles that are internally cold but moving with a large mean velocity in the lab frame. The use of an inhomogeneous field to deflect a molecular beam goes back at least to the experiments of Stern and Gerlach [28] in the 1920s, who demonstrated space-quantisation with the magnetic deflection of a beam of silver atoms. Bennewitz *et al.* described the focussing of a beam of polar molecules with electric fields in the 1950s [29]. Moving beyond lateral focussing to attempt to manipulate the longitudinal velocities of particles in a beam was attempted in several ultimately unsuccessful

experiments in the 1950s and 1960s [30] and not finally demonstrated until the late 1990s, with the Stark decelerator described by Meijer *et al.* [31, 32].

The Stark decelerator operates by creating a series of inhomogeneous electric fields between pairs of electrodes. A molecular beam is sent between these electrodes, which are switched on and off in sequence, providing a continuous ‘hill’ of potential energy up which the particles travel. This method works with polar molecules, *i.e.* those possessing a permanent dipole moment. The alignment of the dipole against the direction of the electric field causes it to be in either a ‘low-field-seeking’ or ‘high-field-seeking’ state, *i.e.* it experiences a force due to the gradient of potential energy in the direction in which the field strength either decreases or increases. As the beam passes between pairs of electrodes they are switched off before the molecules can ‘roll down’ the potential hill on the other side, so that they are constantly decelerated. Meijer *et al.* reported the deceleration of metastable CO [31] and ND<sub>3</sub> [33].

A problem with Stark (and later, Zeeman) decelerators is the undesirable coupling between longitudinal and transverse motion that causes particles to be lost from the manipulated beam. One method to solve or at least ameliorate this problem is to operate the decelerator in a mode such that some of the electrode pairs are used for transverse focussing rather than deceleration, as described by Scharfenberg *et al.* [34]. As an alternative, ‘moving trap’ devices were developed, in which a three-dimensional trapping potential is synthesised and then transported along the beam axis. This was initially demonstrated in a chip-based experiment by Meek *et al.* [35] and in a full-scale decelerator by Osterwalder *et al.* [36]. Travelling-wave Stark decelerators have been used to decelerate, for example, YbF [37] and SrF [38]. Quintero-Pérez *et al.* describe a combined conventional and travelling-wave Stark decelerator used to decelerate NH<sub>3</sub> and ND<sub>3</sub> [39]. Shyur *et al.* reported an improved travelling-wave decelerator [40]. This was later combined with a linear ion trap and time-of-flight mass spectrometer [41]. More details on Stark deceleration can be found in Hogan *et al.* [42] and Meijer [43].

The Zeeman deceleration technique is analogous to Stark deceleration but utilises inhomogeneous magnetic rather than electric fields. The technique works on paramagnetic particles. Since Zeeman deceleration is the main topic of this thesis the subject will be treated in more detail in sections 2.3.2, 2.3.3 and 4.6. The effect of inhomogeneous magnetic fields on paramagnetic particles was demonstrated in the famous Stern Gerlach experiment [28], with these techniques used in a variety of beam-steering experiments [29]. Magnetic traps for neutral particles operate on the same principle as Zeeman decelerators, *i.e.* they confine paramagnetic particles using Stern-Gerlach type forces. Initially proposed in the 1960s [44], they were demonstrated in 1978 with the trapping of neutrons [45]. In 1985 Migdall *et al.* demonstrated the first trapping of a neutral atomic species, with laser-cooled sodium [46]. Although incapable of producing a cooling effect, the magnetic trap may be useful in confining a sample whilst some other cooling effect occurs. Following the advent of Stark decelerators in the late 1990s the first Zeeman decelerators were reported in the mid 2000s [47, 48].

Zeeman decelerators are prone to the same issue previously described for Stark decelerators, namely that coupling between longitudinal and transverse motion causes particles to be lost from the decelerated beam. Again, possible solutions are to operate some of the deceleration coils in focussing mode [49], or the addition of extra focussing elements [50]. Machines with effectively moving

three-dimensional traps were then reported in the early 2010s with two different approaches to coil geometry [51, 52]. The experiment described in this thesis falls into the latter category [1, 2]. A third, radically different field geometry of travelling-wave decelerator has recently been described [53]. Liu and Luo propose to co-decelerate Li with  $\text{CH}_3$  [54] in a moving-trap Zeeman decelerator [55]. The Narevicius group has achieved notable success with their experiment, for example reporting the deceleration of a dual-species beam of molecular oxygen and metastable argon [56], the deceleration and trapping of oxygen and lithium [57, 58], and the deceleration and trapping of oxygen and carbon atoms [59]. A detailed review of magnetic deceleration and trapping can be found in Jansen and Merkt [60].

Reduction of spectral ‘congestion’ may be achieved by eliminating some states from a population. Jansen *et al.* describe exploiting the state-selectivity of Zeeman deceleration to reject certain fine-structure states from a beam of  $\text{He}_2$ , simplifying the identification of components of the spectrum [61]. They were able to further simplify matters by utilising non-adiabatic losses in weak-field regions of the decelerator. Typically these are something to avoid, since they cause magnetically trapped particles to ‘spin-flip’ to an untrapped state and be ejected from a trap [62].

The method of evaporative cooling works by gradually decreasing the depth of the trap confining the sample so that the faster particles escape, with the remainder re-thermalising. Evaporative cooling is an important technique, leading to many significant breakthroughs, such as the first observations of Bose-Einstein condensation [63, 64]. As well as atoms, the method has been applied to molecular species, for example Stuhl *et al.* reported the microwave-forced evaporative cooling of Stark decelerated hydroxyl radicals in a high-gradient magnetic quadrupole trap [65]. More recently Valtolina *et al.* have demonstrated evaporative cooling of  $\text{KRb}$  [66]. The nature of the collisions that occur varies by species and state, and generally it is the case that elastic collisions must dominate over inelastic. The latter can cause molecules to transition to untrapped states and be ejected from the trap.

‘Sympathetic’ or ‘collisional’ cooling refers to a process in which two species come into thermal contact when confined together. The trapped species may be some combination of atoms, molecules, radicals or ions in various types of trap, such as optical [67], radio frequency [68] or magnetic traps [69]. For example, Lim *et al.* considered trapping  $\text{CaF}$  molecules in a microwave trap overlaid with either Li or Rb atoms in a magnetic trap. The atomic species would then be evaporatively cooled, extracting energy from the molecules via thermal contact [70]. Sympathetic cooling of molecules is, however, strongly dependent on the ability to control the types of collision that occur. Inelastic collisions can cause changes in internal state of one or more particle, leading them to be lost from the experiment [71]. Reactive collisions are also undesirable as they alter the species composition of the sample. These pathways need to be suppressed such that only elastic collisions occur. Son *et al.* recently demonstrated cooling of  $\text{NaLi}$  with Na in an optical lattice [72]. Determining the likelihood of success in such an experiment in advance is complex, since small uncertainties in calculated potential energy surfaces can cause large changes in collision cross-sections [73]. Some probabilistic approaches attempt to circumvent the need for precise knowledge of the potential energy surface, for example the work of Morita *et al.* who calculated elastic scattering probabilities for  $\text{CaH}$  and  $\text{SrOH}$  with Li, and  $\text{SrF}$  with Rb [74].

An important technique for atoms since it was demonstrated in the 1980s [75], laser cooling is more difficult to apply to molecular species. The operating principle of the technique is the repeated absorption and emission of photons by the target species. This must occur hundreds or thousands of times for an experimentally useful amount of energy to be removed from the sample, so the transition on which cooling occurs must be ‘closed’, *i.e.* after the an atom is excited by absorption of a photon, it must decay back to its original level [76]. Molecules, however, generally have more complicated energy levels than atoms due to the presence of rotational and vibrational degrees of freedom. This can result in a large number of decay channels, spoiling the closed loop transitions required for laser cooling. Even where laser-cooling is in principle possible for a particular species it may be experimentally very complex to do so, due to the number of different laser wavelengths that may be necessary in order to repump molecules to suitable states. The first laser-cooling of a molecule was not demonstrated until 2010, when Shuman *et al.* reduced the transverse temperature of a beam of SrF to several mK [77]. More molecular species have since been laser-cooled, for example YO [78, 79], CaF [80, 81, 82], YbF [83], with many more under consideration. More information on the laser cooling of molecules can be found in Schreck and van Druten [84] or Softley [9].

Laser-cooling as originally demonstrated was so-called ‘Doppler’ cooling, in which the frequency of light is detuned below the frequency of the cooling transition. This gives a direction-dependence on the cooling effect, since the atoms are only resonant with the laser when they move in an opposing direction to it, Doppler shifting the light. If this is extended to all directions with three pairs of counter-propagating lasers the atoms are slowed in all directions, termed an ‘optical molasses’. Atoms can still drift away from the cooling region, since it is not position-dependent. Adding a quadrupole magnetic field can add trapping as well as cooling. Unlike the purely magnetic trap, the restoring force is not based on Stern-Gerlach effects, but rather the position dependence of light-scattering [85]. Such ‘magneto-optical traps’ or MOTs are a standard tool in atomic physics, and can be extended to molecular species that are laser-coolable, for example SrF [86], CaOH [87, 88], CaF [89, 90], or YO [91]. Dual-species MOTs have existed for some time with atomic species, *e.g.* Rb / Ar [92], Yb / Cs [93] or Na / K [94], and more recent work is extending this to include molecules, *e.g.* Rb / CaF [95]. Some reviews of laser-cooling of molecules can be found in McCarron [96] and Tarbutt [97].

Optoelectrical cooling of polar molecules is an example of a ‘Sisyphus’ technique described by Zepf *et al.* in 2009 [98] and realised for molecules in 2012 using CH<sub>3</sub>F [99]. The molecules are held in an electric field trap with a shallow bottom and steep sides, creating a state- and position-dependent potential. In the centre of the trap the molecules are optically pumped to the more strongly trapped state, and as they move to the outer parts of the trap an RF field drives them to the more weakly trapped state. They therefore have a lower potential and do not regain the energy they lost when moving away from the centre of the trap. This method avoids one of the key difficulties of laser-cooling, namely the requirement for many iterations of the cooling transition, since the energy loss is very high. Compared to the hundreds or thousands of cycles needed in laser cooling only several or even one cycle may be required to reach a target temperature.

Zeeman-Sisyphus deceleration is a related technique, combining a spatially-varying magnetic field with optical pumping. A beam of molecules in a low-field-seeking state climbs up a magnetic ‘hill’, analogous to the operation of a conventional Zeeman decelerator, but rather than the field being

switched off before they roll back down the other side they are optically pumped to a high-field-seeking state so that they are still going uphill. The reverse happens as they reach the next potential maximum. Fitch and Tarbutt described this system in application to CaF [100]. Augenbraun *et al.* demonstrated the deceleration of a beam of CaOH to speeds suitable for magnetic trapping with such a machine [101]. This was later extended to YbOH [102]. A Zeeman-Sisyphus decelerator is currently being built at Durham in the group of H. Williams.

Photoassociation is a method of assembling cold molecules from laser-cooled atoms. Two scattering atoms are driven into an excited bound state by a laser [103]. If the bound state is not to decay back to unbound atoms it must de-excite into a molecular ground state. Homonuclear interactions, in which the two atoms are of the same species, tend to have higher reaction rates than heteronuclear interactions but heteronuclear dimers are more likely to spontaneously decay to a bound ground state [104]. These two factors somewhat cancel, so that rates of both types of association are comparable. All-optical production of Sr<sub>2</sub> has been demonstrated [105, 106], as well as NaCs [107]. Franzen *et al.* recently investigated the possibility of photoassociation of YbRb [108]. As well as the formation of dimer molecules from the photoassociation of two atoms, there has been interest in the possibilities of assembling trimers, for example Rb<sub>3</sub> created from 3 Rb atoms or an Rb<sub>2</sub> molecule and an atom [109]. Another study has considered the association of NaCs with either Na or Cs [110].

A similar approach to photoassociation is magnetoassociation, which utilises Feshbach resonance [111, 112]. In this case an applied magnetic field is used to drive the colliding atoms into a bound state [113]. This is a widely used technique for the creation and study of ultracold molecules, although it is limited to those constituent species that can be laser-cooled (or otherwise brought to the ultracold regime before association). The excited molecules may be driven to ground state by a method such as ‘STIRAP’, stimulated Raman adiabatic passage [114], rather than relying on spontaneous decay. It is typically applied to alkali / alkali earth species. Many species have been studied using Feshbach association, for example NaK [115], NaLi [116, 117], NaCs [118], NaRb [112], RbCs [119]. Work is being done to investigate Feshbach resonances in combinations such as Yb / Li [120], Cs / Yb [121], Rb / Yb [122]. There is also interest in the possibility of forming triatomic molecules via Feshbach resonance, for example with Rb and CaH [123]. Alternative methods to Feshbach association have been demonstrated, such as associating atoms in a dipole trap [124].

Whereas the indirect methods photoassociation and magnetoassociation produce a cold molecular species by assembling it from cooled atoms, the photostop method instead breaks down an existing molecular species [125]. A molecular beam of some precursor species is photodissociated to produce ground-state product species, which move away with some recoil velocity. The laser wavelength and beam velocity are tuned so that the recoil of the daughter species moving counter to the original flow direction cancels out the parent flow velocity, producing particles that are almost stationary in the laboratory frame. The method was proposed in 2007 by Matthews *et al.* [126] with a proof-of-concept experiment reported in 2009 by Zhao *et al.* [127], in which slow NO molecules were produced from NO<sub>2</sub>. This latter experiment was repeated more successfully in an independent result in 2011 described by Trottier *et al.* [128]; in the same year atomic Br was produced from Br<sub>2</sub> via photostop [129]. Rowland reported using photostop to produce SH from H<sub>2</sub>S in 2013 [130]. In 2017 Eardley *et al.* [131] extended this work by trapping SH in a magnetic trap, and also produced O from NO<sub>2</sub> and trapped it [132].

## 1.2 Motivation

Cold atoms and molecules have broad application in precision spectroscopy. Their reduced Doppler widths and cold internal degrees of freedom reduce the uncertainty in measurements. These can be compared with *ab initio* calculations [133]. Spectral resolution is affected by broadening due to Doppler effects and the interaction time of particles with the applied fields that drive transitions of interest. A transition with central frequency  $\nu_0$  is Doppler-broadened such that its full-width half maximum (FWHM)  $\nu_D$  is related to the particle velocity component parallel with the field  $v_{\parallel}$  [134]:

$$\nu_D = 2\nu_0 \sqrt{\ln 2} \frac{v_{\parallel}}{c} \quad (1.2.1)$$

The transit-time FWHM broadening  $\Delta\nu_T$  is also proportional to the particle velocity. Supersonic beams can provide a usefully collision-free environment, and are generally collimated by passage through a skimmer which reduces transverse velocity components, reducing Doppler broadening perpendicular to the beam direction. Depending on the species such a beam may be Stark or Zeeman decelerated to reduce the mean longitudinal velocity. Improved spectral resolution through reduction of transit time of a Stark-decelerated beam was demonstrated by van Veldhoven *et al.* [135] using ND<sub>3</sub> in 2004, and by Hudson *et al.* using OH in 2006 [136]. In 2018, Semeria *et al.* were the first to achieve such a result with Zeeman deceleration, using metastable He<sub>2</sub> [137].

Cold molecules are of interest in the search for new physics. Much work has been done to measure the electric dipole moment of the electron ('eEDM'), in tests of the Standard Model. Various models predict different values that this moment should take, and experiments have been performed to place constraints on its value. Typically these experiments search for energy shifts arising from the interaction of the eEDM with external electric fields, an effect which is enhanced by performing the measurement on a heavy and paramagnetic species [138]. Kara *et al.* used a supersonic expansion of YbF [139], later improved upon by Ho *et al.* [140]. Cairncross *et al.* and then Roussy *et al.* used HfF<sup>+</sup> molecular ions [141, 142].

At low temperatures the effects of external fields can dominate the collision energy of molecules, allowing the possibility of controlling collisions. Many chemical reactions are inhibited by an activation energy barrier and so at low temperatures they are classically forbidden. Cold molecule experiments may allow study of the quantum mechanical nature of chemical reactions, for example in crossed-beams [143]. Other non-classical effects such as resonances or reflections may be studied in, for example, merged beam experiments [24]. Studies of gas-phase molecular energy transfer and chemical reactions at very low temperatures can be applied in modelling interstellar molecular synthesis [144]. Deceleration can produce low-velocity beams with relatively few quantum states present, allowing more precise study of reactive, elastic or inelastic collision processes, for example in crossed beams of xenon with Stark-decelerated OH [145]. For a recent review of the prospects of ultracold chemistry see Heazlewood and Softley [4].

Another area of study that utilised cold atoms and molecules is quantum simulation. This field arguably originated in 1982 when R. Feynman suggested that using classical computation to study quantum phenomena was not entirely feasible [146]. This arises from the vast number of long-range interactions between all the particles in a many-body system. Bulk effects such as high-temperature superconductivity are thus almost impossible to understand precisely. Feynman proposed using systems of interacting quantum particles prepared in specific states with tunable, shorter-ranged interactions. Rather than attempting to generate and solve a specific Hamiltonian classically, an analogous Hamiltonian could then be engineered on this simulator and used to infer information about the ‘real’ problem system [147]. Molecules can have much more complex energy levels than atoms, as mentioned in the context of them being difficult to laser cool, but for quantum simulation this offers an opportunity since they can potentially be used to model a wider variety of physical systems [148, 149]. Recent proposals suggest that molecular spin ‘qudits’ offer a powerful technique to create quantum simulators [150].

### 1.3 Aims Of This Project And Thesis

As stated earlier, the initial research goal of the decelerator project was related to quantum simulation, namely the production of a cold beam of a species such as CaF that could be trapped and further cooled in a molecular magneto-optical trap. The focus shifted towards cold chemistry and a study of collision dynamics with species such as Li and CaH, which could be decelerated and loaded into a proposed hybrid trap, the MT-MOT (see chapter 5). More recently, we have been working towards applications in the search for new physics via precision spectroscopy, specifically considering deceleration and sympathetic cooling of atomic H using Li.

Key to these goals is proving the viability of the travelling-wave Zeeman decelerator that we have constructed, characterising its performance and making predictions of its capabilities, both as it is and with proposed upgrades. Previously, we did not have a good understanding of the data being produced by the experiment, because the multiple-trap design by necessity produces a complex output signal, with multiple flight profiles superimposed. This thesis concerns the development of simulation methods to allow us to untangle these results and provide a good quantitative, as well as qualitative, description of our data. These methods have been fed back in to selecting operational parameters for the decelerator, allowing us to take better data. Chapter 2 gives a review of the decelerator design and operational theory. This includes the source, deceleration stage, and detector. Details of how these components are simulated is given in detail in chapter 3, which also applies these methods to real experimental data, namely the guiding and deceleration of metastable argon. Chapter 4 gives some detailed predictions regarding potential performance of the decelerator in the case of argon, which our simulation codes have facilitated. Also included are analyses of various effects in the experiment that were not previously well understood.

Chapter 5 gives some outlook for the experiment, including some modifications to the decelerator. Predictions related to future work, specifically the deceleration of species relevant to the goals of cold chemistry and the search for new physics, are given. The proposed MT-MOT is described, briefly. This was studied in previous work by the current author [151], and subsequently by other members of the group using some of the computer codes developed and described in this thesis. Some results related to this subject are also presented here.

## Chapter 2

# Decelerator Experimental Theory and Design

This chapter will discuss the components of the Travelling-Wave Zeeman Decelerator, including pulsed-source Even-Lavie valve, deceleration coils, transverse focussing quadrupole guide, and microchannel-plate detector, as shown in the schematic in figure 2.1. For further details on the vacuum system, the Even-Lavie valve and the General valve that it replaced see Mizouri [1]. For further details of the deceleration coils, electromagnetic quadrupole and power electronics see McArd [2].

### 2.1 Experimental Layout

The source chamber contains the Even-Lavie valve, which has a cooling jacket fed by a total-loss liquid nitrogen system. At right angles to the valve nozzle is a filament for injecting thermal electrons into the source to help stabilise the dielectric-barrier discharge when producing metastable argon; this is powered by a small constant-current DC supply running at 3 V with a typical current of 2 A. Also present is a fast-ionisation gauge (FIG) on a swivel which can be rotated into the path of the molecular beam as it exits the skimmer. A turbopump backed by a rotary pump evacuates the chamber.

Exiting the source chamber is a stainless steel tube through which the sample passes to the detection chamber. The deceleration coils, formed as pairs of flattened helices, are located above and below the tube (outside of vacuum). There are currently 4 modules, each connected as 4 individual phases. The overall length of the deceleration stage is  $\approx 0.5$  m, and the system is designed in a modular fashion so that extra modules could be added. Machined cooling blocks with feeds for a closed-loop chilled water supply sit above and below the coils. An additional quadrupole guide provides transverse focussing; initially this was formed of eight long wire segments in one circuit of approximately 600 A at 37 V DC. This was removed, and is due to be replaced by a series of custom-manufactured permanent magnets.

The detection chamber, also evacuated by a turbopump / rotary backing pump, currently contains a microchannel-plate (MCP) detector, which is connected to an oscilloscope via a bespoke impedance-matching buffer amplifier. This was used for characterising the detector with metastable argon, and can be swapped out for a quadrupole mass spectrometer (QMS) when working with other species.

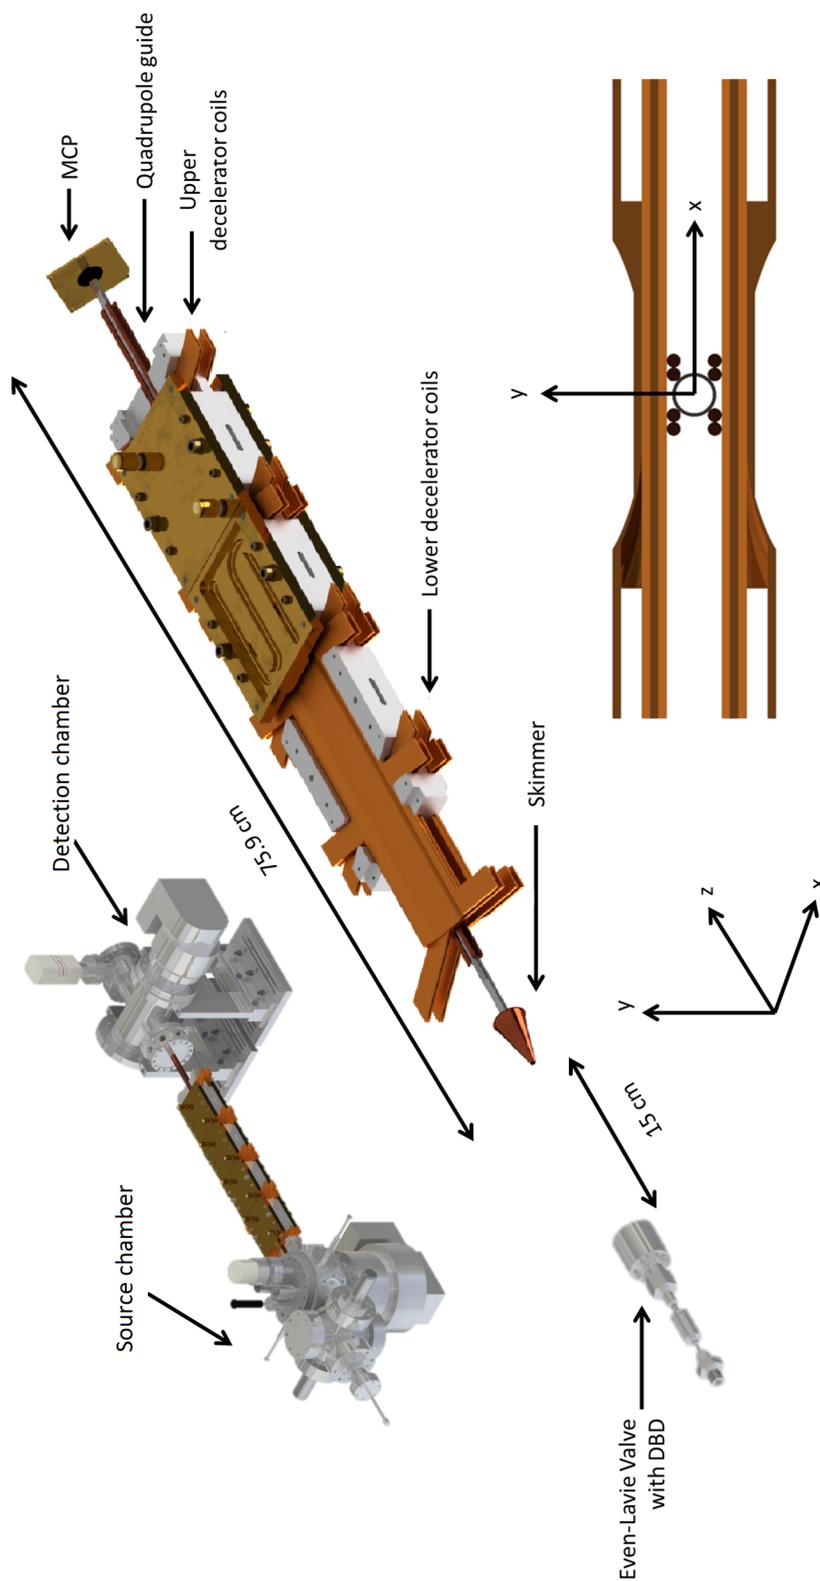


Figure 2.1: Schematic of the Travelling-Wave Zeeman Decelerator. Shown are the differentially-pumped source and detection chamber, the Even-Lavie valve with dielectric-barrier discharge, paired twin flattened helical decelerator coils, quadrupole guide and microchannel-plate detector. Image reproduced from L. McArd [2].

## 2.2 Beam Source

This section gives a brief overview of the theory of supersonic expansions, and the apparatus we use to generate such an expansion.

### 2.2.1 Supersonic Expansion

We utilise a pulsed supersonic expansion source based on a cryogenically cooled Even-Lavie valve. In the context of molecular beams produced by expansion into a vacuum we can consider the two limiting cases defined by the relationship between the mean free path  $\lambda_0$  (or distance between collisions) of the constituent molecules, and the diameter  $D$  of the nozzle of the valve used to produce the expansion. For  $\lambda_0 \gg D$  there is little redistribution of particle energies during expansion, such expansions are termed ‘effusive beams’. At the opposite end of the scale ( $\lambda_0 \ll D$ ), a ‘supersonic expansion’ [18] essentially rearranges the distribution of molecular energies to produce a packet of particles that are internally cooler at the expense of moving at a high central velocity in the laboratory frame of reference. The following description, whilst strictly only applicable to an idealised continuous expansion, captures some of the essential features of the system that are relevant to our goals. Starting with a sample in a stagnation chamber at a temperature  $T_S$  and a pressure of several bar, the distribution of molecular speeds  $v$  of the sample can be described by a Maxwell-Boltzmann speed distribution:

$$f_S(v) = 4\pi \left( \frac{1}{2\pi\sigma_S^2} \right)^{3/2} v^2 \exp\left( \frac{-v^2}{2\sigma_S^2} \right). \quad (2.2.1)$$

This is a Gaussian with variance  $\sigma_S^2$ , defined in terms of the Boltzmann constant  $k_B$ , and the mass  $m$  and temperature  $T_S$  of the molecules in the sample:

$$\sigma_S^2 = \frac{k_B T_S}{m}. \quad (2.2.2)$$

The valve opens to allow an isentropic expansion into the source chamber. The escaping molecules are initially dense enough for rapid re-thermalisation, but the chamber must be at a sufficiently low pressure; shockwaves diminish the expansion, and in order for the isentropic condition to be valid collisions between particles in the beam must dominate over collisions with the chamber background. Early implementations of supersonic expansions were continuous beams and suffered these problems [152]. Utilising a pulsed source improves the situation since the valve repetition rate can be set to minimise the background pressure. A ‘clean’ unhindered expansion is also helped by the valve being open for a time shorter than the typical time taken for a particle from the beam to reflect from the chamber wall [153]; the short opening times ( $\approx 20 - 30 \mu\text{s}$ ) achievable with an Even-Lavie valve facilitate this. Our system is differentially pumped, with the source chamber held at a background pressure of order of  $10^{-7}$  mbar. When working with argon a typical repetition rate of 2 or 3 Hz is used, during which the pressure in the source chamber rises by an order of magnitude. The valve and pumps are capable of a higher repetition rate without adversely effecting the expansion, but we utilise a dielectric barrier discharge (DBD) incorporated into the valve body to energise argon to a metastable state, along with a thermal-emission filament to inject electrons into the beam to help stabilise the discharge. Both of these cause heating in the valve, and keeping the repetition rate low helps mitigate this. As will be shown, it is desirable to produce as cold a source as possible in order to achieve slower, internally colder beams. A skimmer cores the beam immediately post-expansion,

removing the outer parts of the packet for which the adiabatic assumption is least valid. As the packet expands the density drops until the rate of collisions essentially reaches zero, producing a ‘terminal’ speed distribution of the form [154]:

$$f_T(v) = 4\pi \left( \frac{1}{2\pi\sigma_T^2} \right)^{3/2} v^2 \exp\left( -\frac{(v - v_T)^2}{2\sigma_T^2} \right). \quad (2.2.3)$$

The packet of molecules is now moving in the laboratory frame with a mean velocity  $v_T$  and a narrower variance  $\sigma_T^2$ , since the terminal translational temperature  $T_T < T_S$ :

$$\sigma_T^2 = \frac{k_B T_T}{m}. \quad (2.2.4)$$

This is making a simplification that the radial and longitudinal temperatures are approximately equal; this is not too important since the skimmer constrains the range of possible transverse velocities. It would be more accurate to describe the supersonic expansion in terms of a velocity distribution [155] in which the particles now have some distribution of radial / longitudinal velocity components  $v_r / v_z$  with variance  $\sigma_{T,r}^2 / \sigma_{T,z}^2$  (and therefore temperatures  $T_{T,r} / T_{T,z}$ ):

$$f_T(\mathbf{v}) = \frac{1}{2\pi\sigma_{T,r}^2} \left( \frac{1}{2\pi\sigma_{T,z}^2} \right)^{1/2} \exp\left( -\frac{v_r^2}{2\sigma_{T,r}^2} - \frac{(v_z - v_T)^2}{2\sigma_{T,z}^2} \right). \quad (2.2.5)$$

For our purposes, however, we can use the terminal temperature  $T_T$ . The average total molar energy  $\langle E \rangle$  is given by the sum of the enthalpy  $H$  and the kinetic energy  $E_k$  due to the bulk motion of the atoms. Enthalpy is defined as  $H \equiv U + PV$ , the internal energy  $U$  plus the work done at pressure  $P$  to change the molar volume  $V$ . The total energy  $\langle E \rangle$  is conserved so the states in the source chamber and in the supersonic expansion can be compared:

$$H_S + E_{k,S} = H_T + E_{k,T}. \quad (2.2.6)$$

The enthalpy change  $\Delta H = H_S - H_T$  can be shown to be equal to the kinetic energy of the terminal bulk motion, since the bulk motion in the stagnation chamber is zero:

$$\Delta H = E_{k,T}. \quad (2.2.7)$$

The kinetic energy  $E_{k,T} = \frac{N_A m v_T^2}{2}$ , where  $m$  is the particle mass and  $N_A$  is the Avogadro number, giving the velocity  $v_T$  in terms of  $\Delta H$ :

$$v_T = \left( \frac{2\Delta H}{N_A m} \right)^{1/2}. \quad (2.2.8)$$

The temperature  $T$  can be related to the enthalpy  $H$  via the relation  $T = c_p H$ , obtained by assuming that the isobaric heat capacity  $c_p = (\partial H / \partial T)_p$  has negligible dependence on temperature or pressure, so:

$$v_T = \left( \frac{2c_p \Delta T}{N_A m} \right)^{1/2}. \quad (2.2.9)$$

Equation 2.2.9 is of limited applicability in the context of a supersonic expansion, in which  $\Delta T$  is typically large; for example, when using argon in our Even-Lavie valve we observe  $\Delta T \approx 130$  K, over which range  $c_p$  varies significantly. The approximation can be improved by including the

constant volume heat capacity  $c_v$ , which is related to  $c_p$  by  $c_p = c_v + R$  and  $\gamma = c_p/c_v$ , where  $R$  is the gas constant, and also considering  $T_T \ll T_S$ :

$$v_T = \left( \frac{2R\gamma T_S}{(\gamma - 1)N_A m} \right)^{1/2}. \quad (2.2.10)$$

The ratio  $\gamma$  takes the value of 5/3 for ideal monatomic species and 7/5 for ideal diatomic species. Variations in  $c_p$  and  $c_v$  when considering a real system partially cancel each other out.

The Mach number  $M$  is the ratio of gas flow speed to the local speed of sound  $c$ , with:

$$c = \left( \frac{\gamma k_B T}{m} \right)^{1/2}. \quad (2.2.11)$$

By definition, in a supersonic expansion the Mach number is greater than unity. Various factors determine the value of  $M$  for a given expansion, including the dimensions of the valve, the pressure and temperature of the source, the species involved, and the cross-sections of the various possible collisions, whether elastic or inelastic. The temperatures  $T_S$  and  $T_T$  can be related using the Mach number [19]:

$$T_T = T_S \left( 1 + \frac{\gamma - 1}{2} M^2 \right)^{-1}. \quad (2.2.12)$$

Equations 2.2.10 and 2.2.12 illustrate the dependence of  $v_T$  and  $T_T$  on  $T_S$ , and hence the importance of a cold source in order to make slow, internally cold beams.

## 2.2.2 Even-Lavie Valve

We utilise an Even-Lavie pulsed valve [156] in order to produce a supersonic expansion beam. Some background on supersonic beams is given in section 2.2.1. The valve incorporates a dielectric-barrier discharge (DBD) unit which can be used to excite certain species into metastable states, see section 2.2.3.

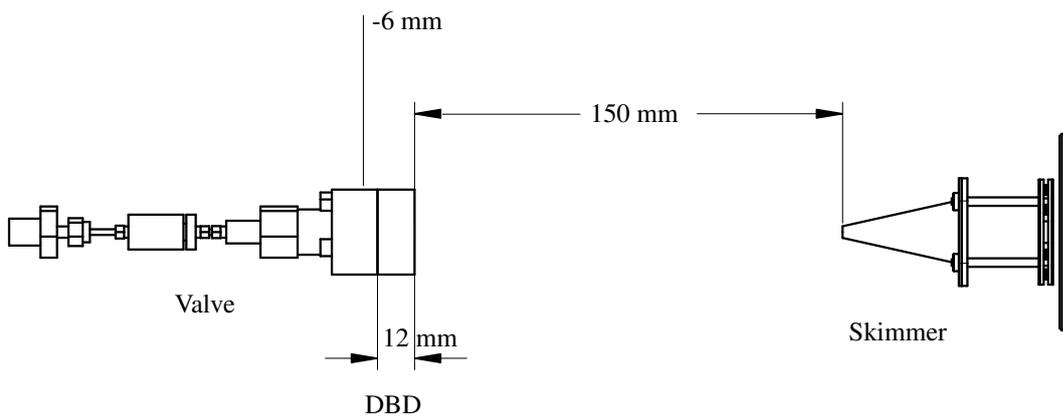


Figure 2.2: Internals of source chamber, Even-Lavie valve with dielectric-barrier discharge and skimmer. Defining the origin of the beam propagation axis at the front of the valve housing, the valve nozzle is at  $-6$  mm.

The experiment initially utilised a General Valve with a glow discharge system [1] but later moved to an Even-Lavie valve [2]. The latter produces a high-density, short duration gas pulse due to hav-

ing an opening time of order  $\approx 30 \mu\text{s}$ . The valve is cooled via a total-loss liquid nitrogen system, typically reaching temperatures of order  $\approx 140 \text{ K}$ . Working with  $^3\text{P}_2$  metastable argon at a backing pressure of 6 bar produces a beam with a translational temperature  $\approx 6 - 8 \text{ K}$  moving in the laboratory frame with a flow velocity of  $\approx 350 - 400 \text{ m s}^{-1}$ ; see sections 2.4.1 and 3.2 for details of how these parameters are extracted from experimental data.

Figure 2.2 shows a schematic of the valve assembly (see also figure 2.1). The rear of the valve is at the left, which is where the gas feed line is connected. The cooling jacket (not depicted) surrounds the main body at the front. Also not depicted are the electrical connections for the valve actuator, DBD and thermocouple temperature probes. We define the beam propagation direction along the  $z$ -axis from left to right, with the origin at the front face of the valve. The emission nozzle is recessed so that the atoms enter the beam at  $\approx -6 \text{ mm}$ . The DBD is 12 mm in length; as described in section 3.2 the supersonic expansion should have completed by the time the gas packet passes through this. The atoms pass to the skimmer, which is located 150 mm downstream and has a 4 mm diameter; this serves to ‘core’ the beam, removing atoms with high transverse velocities. Skimmer placement can be very important in experimental setups such as this, so as not to adversely affect the beam expansion [157], and the layout shown has been found to work well for our apparatus [2].

Our group previously reported [2] estimated flow velocities  $\approx 330 - 370 \text{ m s}^{-1}$  and translational temperatures of  $\approx 3 - 4.5 \text{ K}$  for argon, compared to values of  $\approx 320 \text{ m s}^{-1}$  and  $\approx 1.5 \text{ K}$  given in Even [157]. Firstly, it is worth noting that Even was working at stagnation pressures of 50 bar and valve temperatures of 130 K, and that the numbers quoted were estimated assuming a pure adiabatic expansion. Secondly, our previously published figures are almost certainly under-estimates; we have subsequently developed more accurate models for fitting our data, as discussed in sections 3.1.1 and 3.2. Characterising molecular beam experiments such as this can be difficult, and dependent on the dynamical model used.

### 2.2.3 Dielectric Barrier Discharge

Molecular beam experiments often employ some mechanism to excite atoms or molecules, producing ions or radicals. When working with argon this is necessary for two reasons; firstly, in order for the decelerator fields to manipulate the atoms they must be in a paramagnetic state, and secondly in their ground-state they do not possess sufficient energy to trigger the microchannel-plate detector (section 2.4.2) Detection would therefore require an alternative technique such as multi-photon ionisation. This would also be needed for proposed hydrogen experiments. Earlier versions of our experiment utilised electric discharge excitation [1], in which an applied potential between electrodes ionises atoms to produce a mixture of atoms and plasma. In this environment some atoms are excited to useful metastable states via collisions. We now utilise a dielectric-barrier discharge (DBD) which is attached to our Even-Lavie valve. The DBD is formed of concentric electrodes with a dielectric sandwiched between them, and the beam passes through the centre. This produces a more uniform excitation with less heating [158]. Since the flow velocity of a supersonic expansion correlates with the source temperature this is beneficial. The discharge is initiated by applying a series of radio-frequency pulses to the outer electrode, with parameters such as the number of pulses and frequency controlled via a software interface.

### 2.2.4 Stability of the Source

When characterising the decelerator with metastable argon produced in the Even-Lavie valve / dielectric-barrier discharge, we experienced problems with fluctuations of the time-of-flight signal received at the microchannel-plate detector. We were able to largely eliminate the effect of these fluctuations from our data by taking each data set as an average over 40 individual signals, with the averaging performed directly on the oscilloscope before the data were saved to file. The pulse rate of the decelerator is variable, and we typically operate it at 2 Hz to reduce the heat load in the valve arising from the operation of the discharge. Running 40 sweeps of each set of parameters is not a significant experimental handicap, since reconfiguring the field-programmable gate array (FPGA) controller for the decelerator power electronics for a new set of parameters takes several minutes anyway.

Fluctuations in the signal may have several root causes such as an excessively high background pressure interfering with the gas pulse expansion, or ringing on the detector circuit caused by an impedance mismatch between the detector and the oscilloscope (see section 2.4.3). More likely is instability in the excitation scheme. In the first iteration of our Zeeman decelerator a General valve was used [1], this was fitted with a thermionic-emission filament [159] to assist with the glow discharge mechanism. Such a system was not originally installed with the Even-Lavie valve but was added to see if it would improve the shot-to-shot stability. The unit used is an Agar Scientific AGA092 which is mounted inside the source chamber, off the beam axis and pointing toward the valve assembly. The filament is connected to a variable output power supply that operates at 2 V and 0 – 3 A.

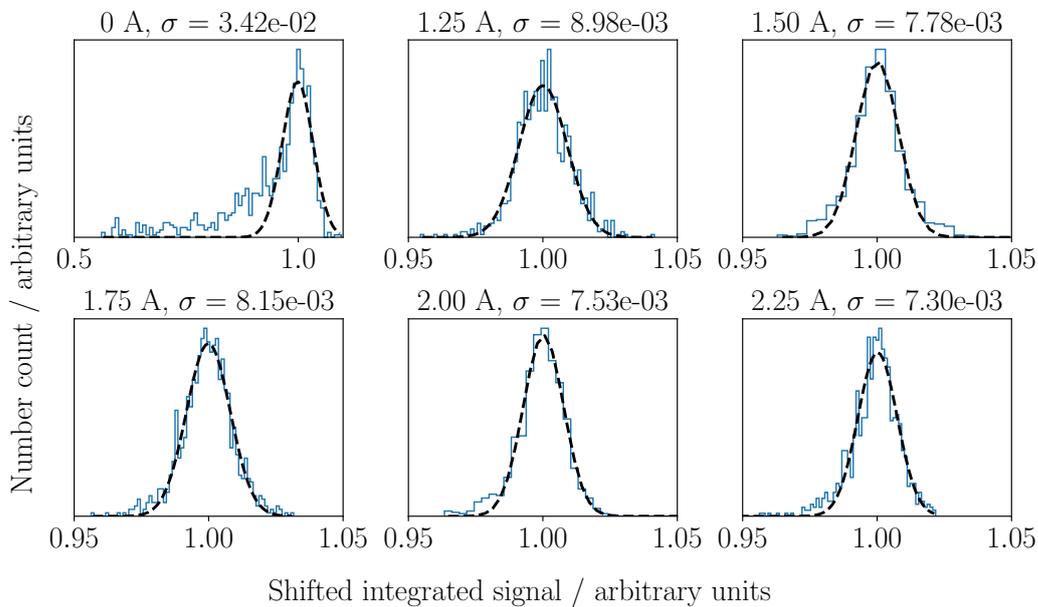


Figure 2.3: Effects of the electron-emission filament on shot-to-shot stability for metastable argon, for different values of the filament current. For each plot 1000 signals are captured, their integrated areas are histogrammed and then fitted with a Gaussian to extract the standard deviation  $\sigma$ . Currents of  $\geq 1.25$  A give a factor of  $\approx 4$  increase in stability.

In order to quantify the effects of the filament we ran the beam with the electronics disengaged

such that the beam passed through the decelerator to the MCP with no manipulation. No cooling was applied to the valve. Disabling the averaging function on the oscilloscope we collected 200 sets of data with the filament switched off, then 200 with it activated for each of several different current values. The whole process was repeated 5 times to give 1000 samples for each value of the current; the data files were trimmed to only include the points bracketing the signal peak, and were then numerically integrated using Python. When our source has been operating for several hours the beam properties drift; in order to remove any effects arising from this each block of 200 data points was normalised against the mean value for the block. The data were combined and fitted with Gaussian functions using least-squares fitting to extract values of the standard deviation. The results are shown in figure 2.3, which shows that even a moderate current through the filament improves the stability of the discharge by a factor of approximately four. There is little point in driving the filament at its highest output, and it is in any case worth running at as low a current as possible for two reasons; to increase the lifespan of the filament, and to reduce the heat load into the valve. Even at  $\approx 1.25$  A there is a noticeable effect on the temperature of the source over time.

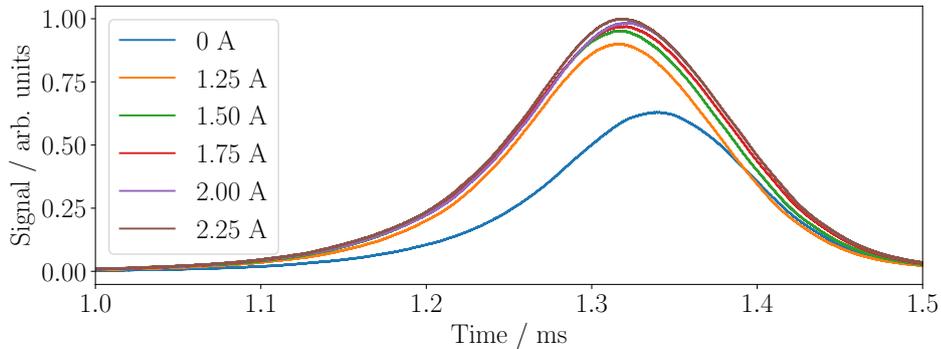


Figure 2.4: Effects of the electron-emission filament on signal intensity and arrival time for metastable argon, for different values of the filament current. The electron injection greatly increases the number of atoms detected, with the signal peak increased by almost 50 %, but at the expense of also increasing the beam flow velocity. See section 2.2.1 for the effects of increased source temperature on supersonic expansions.

This analysis tells us nothing about the effects of electron injection on the beam temperature and intensity. To quantify this, we took a block of 200 time-of-flight data samples for each of the listed values of the filament current and averaged them, and they can be compared as in figure 2.4. A significant increase in the number of metastable argon atoms triggering the detector can be seen for all values  $\geq 1.25$  A; however, there is also a decrease in arrival time and hence increase in beam velocity. Fitting the data with the model given in section 3.1.1 gives an estimated beam velocity of  $693 \text{ m s}^{-1}$  and translational temperature of  $5.2 \text{ K}$  when the filament is off; for the data with  $1.25 \text{ A}$  these values increase to  $703 \text{ m s}^{-1}$  and  $5.6 \text{ K}$ . These tests were conducted with the valve at room temperature, whereas under normal operation the cryogenic cooling system would be active. With the valve cooled we have not repeated this work in as much detail, but comparisons of a few sets of data taken with and without the filament active suggested that for a cooled beam the effects on the beam flow velocity and translational temperature are negligible, and compensated for by the increase in stability and amplitude of the detected beam. The values given for argon in section 2.2.2 ( $\approx 6 - 8 \text{ K} / \approx 350 - 400 \text{ m s}^{-1}$ ) relate to the beam with the filament active. All experimental data utilising argon presented in this thesis were taken with the filament active at  $1.2 \text{ A}$ .

## 2.3 Zeeman Deceleration

### 2.3.1 The Zeeman Effect

Shifts in the energy levels of an atom or molecule in the presence of an external magnetic field are key to a broad range of experiments involving trapping or deceleration of said particles. A brief overview of the principles of these phenomena is given in this section. The incorporation of these effects into our computer simulations is described in section 3.4.

The field of magneto-optics, or the study of the effects of magnetic fields on light, stretches back to at least the 1840s and the work of Faraday. As well as his well-known discovery of the effect a strong magnet had on the polarization of light, in 1862 Faraday conducted an experiment in which he searched unsuccessfully for an effect induced in the spectrum of a flame by the application of a magnetic field [160]. Some decades later Zeeman read of the null result of Faraday and decided to repeat the experiment [161], and in 1896 he showed that the D-lines in sodium spectra were broadened in the presence of a magnetic field [162].

In response to the work of Zeeman, Lorentz predicted that the broadening was actually a splitting of the lines into doublets or triplets, depending on the orientation of the viewing apparatus with respect to the applied field. Although the structure of atoms was at the time unknown, Lorentz made the assumption that they contained some form of oscillating charged particle. In a magnetic field these charges would experience a force acting transverse to their motion (the ‘Lorentz force’); depending on the orientation of the atom with regard to the field this gives two additional rotational frequencies, one smaller and one larger than the frequency in the absence of the external field. Parallel to the field, this gives a doublet of spectral lines with opposite circular polarisation. Perpendicular to the field a triplet is seen, one line with the unperturbed frequency, and one each above and below it, all linearly polarised [85]. Zeeman soon confirmed these predictions in experiments with cadmium, thus also providing early evidence for the existence of the electron (which was explicitly detected the following year by J. J. Thomson).

The explanation of Lorentz was soon shown to be incomplete, however, when further experiments were conducted by Preston in 1897 [163, 164]. Preston found higher order splitting of spectral features which could not be accounted for classically. This ‘anomalous’ (as opposed to the original or ‘normal’) Zeeman effect defied ready explanation for years, with the required pieces of the puzzle falling into place only over the next three decades as the first components of quantum theory were developed. An atom with a magnetic dipole moment  $\boldsymbol{\mu}$  experiences an energy perturbation in an external field  $\mathbf{B}$  dependent on the orientation between  $\boldsymbol{\mu}$  and  $\mathbf{B}$ :

$$U = -\boldsymbol{\mu} \cdot \mathbf{B}. \quad (2.3.1)$$

Departing from the classical view, in what came to be called the ‘old’ quantum theory Bohr developed a model of the atom as having only discrete allowed levels or ‘orbitals’ that electrons could occupy. This explained the discrete nature of spectroscopic lines, but also predicted that the orbital angular momentum of the electrons in the atom is also quantised along the measurement axis, and so the atomic moment can only take certain fixed values with regard to the direction of the applied field. This was not enough to complete the picture until the concept of ‘spin’ was introduced. In 1921 Compton was the first to suggest that electrons possessed some extra form of angular momentum that came to be called spin [165], although it has no direct analogue in the classical world. The idea was developed by Pauli, who in 1925 described how such a two-valued parameter would resolve the problem of the anomalous Zeeman effect [166]; the external field couples to the orbital and spin momenta of the electrons, producing the extra energy level splitting that Preston had discovered. Additionally, coupling between the orbital and spin moments of the electrons with each other and with the spin (if present) of the nucleus contribute to the ‘fine’ and ‘hyperfine’ structure which had also puzzled spectroscopists. The total moment of the atom is a combination of several components. Beginning with the moment  $\mu_L$  arising from the orbital angular momentum  $\mathbf{L}$ , with a g-factor  $g_L$ :

$$\hat{\mu}_L = g_L \frac{-e}{2m_e} \hat{\mathbf{L}} = -g_L \mu_B \frac{\hat{\mathbf{L}}}{\hbar}, \quad (2.3.2)$$

where  $e$  and  $m_e$  are the electron charge and mass, and  $\mu_B$  is the Bohr magneton  $\mu_B = \frac{e\hbar}{2m_e}$ . The spin  $\mathbf{S}$  likewise produces a moment  $\mu_S$  with spin g-factor  $g_S$ :

$$\hat{\mu}_S = g_S \frac{-e}{2m_e} \hat{\mathbf{S}} = -g_S \mu_B \frac{\hat{\mathbf{S}}}{\hbar}. \quad (2.3.3)$$

Lastly, the nuclear spin  $\mathbf{I}$  has moment  $\mu_I$ :

$$\hat{\mu}_I = g_I \frac{e}{2m_p} \hat{\mathbf{I}} = g_I \mu_N \frac{\hat{\mathbf{I}}}{\hbar}, \quad (2.3.4)$$

with g-factor  $g_I$  and nuclear magneton  $\mu_N = e\hbar/2m_p$  where  $m_p$  is the proton mass. The external field couples with these three angular momenta giving the Zeeman energy shift which can be written as the Hamiltonian operator:

$$\hat{H}_{Zee} = \frac{\mu_B}{\hbar} \left( g_L \hat{\mathbf{L}} + g_S \hat{\mathbf{S}} \right) \cdot \mathbf{B} - \frac{\mu_N}{\hbar} \hat{\mathbf{I}} \cdot \mathbf{B}. \quad (2.3.5)$$

This can be simplified by noting that  $g_L = 1$  and  $g_S \approx 2$ . The nuclear magneton is smaller than the Bohr magneton by  $m_p/m_e$ , *i.e.* three orders of magnitude, so that the last term may often be neglected. Defining the space quantisation axis along which  $\mathbf{B}$  is aligned as being in the direction of  $z$  so that  $\mathbf{B}$  has magnitude  $B_z$ , we obtain:

$$\hat{H}_{Zee} = \frac{\mu_B B_z}{\hbar} \left( \hat{L}_z + 2\hat{S}_z \right). \quad (2.3.6)$$

The normal Zeeman effect is the limiting case when the net electron spin is zero, with the anomalous effect being more general; however the historic nomenclature has remained in use. The magnetic moments associated with the spin and orbital angular momenta interact, and if an applied field has an effect smaller than this interaction then  $\mathbf{L}$  and  $\mathbf{S}$  are not individually conserved, and instead effectively precess around the total angular momentum  $\mathbf{J} \equiv \mathbf{L} + \mathbf{S}$ . Under this ‘ $LS$ -coupling’ the atom has an effective magnetic moment  $\mu_J$  and  $g$ -factor  $g_J$ , and the Zeeman Hamiltonian can be written:

$$\hat{H}_{\text{Zee}} = \frac{\mu_B g_J \hat{J}_z B_z}{\hbar}. \quad (2.3.7)$$

$LS$ -coupling breaks down in fields strong enough that the energy splitting is comparable to the spin-orbit interaction, and energy shifts are determined by the ‘Paschen-Back’ rather than Zeeman effect. No further consideration of the Paschen-Back effect will be given, since the field strengths required for this decoupling are typically large for all but the lightest species. For example, much of the work in this thesis involves  ${}^3P_2$  argon in which the spin-orbit coupling gives rise to an energy splitting of  $\approx 0.1$  eV; to reach an equivalent level according to equation 2.3.8 would require a magnetic field strength of  $\approx 1000$  T, compared to  $\approx 0.5$  T in our decelerator.

Methods of solving the Hamiltonian will be discussed in section 3.4; for now we will simply consider some results, using ground state silver as an example. In the  ${}^2S_{1/2}$  term  $\mathbf{L} = 0$  so  $\mathbf{J} = \mathbf{S}$  but retaining  $\mathbf{J}$  in 2.3.7 makes it more general. Silver has nuclear spin  $\mathbf{I} = 1/2$ , this gives four sub-levels which in a moderate strength magnetic field can be labelled by the projection numbers  $m_I$  and  $m_J$ , both taking values  $\pm 1/2$ . Neglecting the interaction of the nuclear spin with the external field, as described above, the Zeeman shift is approximately given by:

$$\Delta E \approx g_J \mu_B m_J B_z, \quad (2.3.8)$$

as shown in figure 2.5. This is known as the ‘hyperfine Paschen-Back’ regime, commonly used in spectroscopy. If the field is inhomogeneous, this gives rise to a force:

$$F_z = -\frac{d\Delta E}{dz} = -g_J \mu_B m_J \frac{dB_z}{dz}. \quad (2.3.9)$$

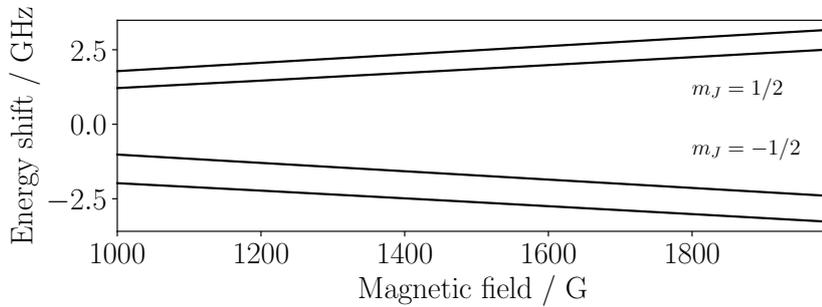


Figure 2.5: Zeeman splitting of ground-state silver as a function of applied magnetic field. Dependent on their projection quantum number the sub-levels will experience force in an inhomogeneous field due to the gradient of potential energy; these are often called ‘low-field seeking’ or ‘high-field seeking’. This is the basis of the Stern-Gerlach effect, the operating principle of Zeeman decelerators.

This was the basis of the 1922 experiment of Stern and Gerlach [28], who used inhomogeneous magnetic fields produced by shaped magnets to deflect beams of silver atoms. In modern times this work is often presented as the archetypal proof of electron spin, but at the time it was conducted spin was not known. The experiment was designed to look for quantisation of orbital angular momentum (in fact, Stern was hoping for a negative result to disprove quantum mechanics, about which he was sceptical). The interpretation of the results was that quantised orbital angular momentum had indeed been found; this was later shown to be incorrect since ground-state silver atoms have  $\mathbf{L} = 0$ . Nonetheless, a clear two-valued splitting of the beam, rather than a classical result, was demonstrated, later understood as ‘space quantisation’ due to spin. The Stern-Gerlach effect is the key operating principle of the Zeeman decelerator and related magnetic traps for neutral particles, which use field minima to confine weak-field seekers.

We were able to discard the nuclear term in the Hamiltonian shown in equation 2.3.5 because the coupling of the nuclear moment to the external field is much weaker than that of the electronic moment, so that equation 2.3.7 does not depend on  $I$ . However, in a more complete view we must include the ‘hyperfine’ interaction of the nuclear and electronic spin which affects the expectation value of the energy, as seen in figure 2.5. This will be discussed in section 3.4, which considers methods of calculating the required energy shifts in Monte Carlo simulations.

### 2.3.2 Zeeman Decelerators

Section 2.3.1 described the Zeeman effect and the related Stern-Gerlach mechanism, by which paramagnetic particles experience a force in an inhomogeneous magnetic field. This is illustrated in figure 2.5, in which it can be seen that ground-state silver atoms may be in one of four sub-levels which minimise their energy in regions of low or high field strength, depending on the projection of their angular momentum with respect to the field. Proposals to utilise this effect for confinement of neutral particles were made in the early 1960s [44] and achieved in the 1980s [46] in so-called ‘magnetic bottles’ or traps. The ‘Zeeman decelerator’ was described by the groups of Merkt [47] and Narevicius [48] in 2007; also based on the Stern-Gerlach effect, the Zeeman decelerator extracts kinetic energy from a molecular beam of particles with a magnetic dipole moment by utilising a series of switched solenoids to create a time-dependent magnetic potential. Low-field-seeking (LFS) particles (*e.g.* those shown in the upper levels of figure 2.5) lose energy and are decelerated as they approach an active solenoid; before they pass the centre point the coil is deactivated. This process is repeated, so that the particles travel up a succession of potential ‘hills’, and the machine may be dozens of sections long. Machines of this type will be referred to as ‘conventional’ Zeeman decelerators, to distinguish them from the travelling-wave or moving-trap type.

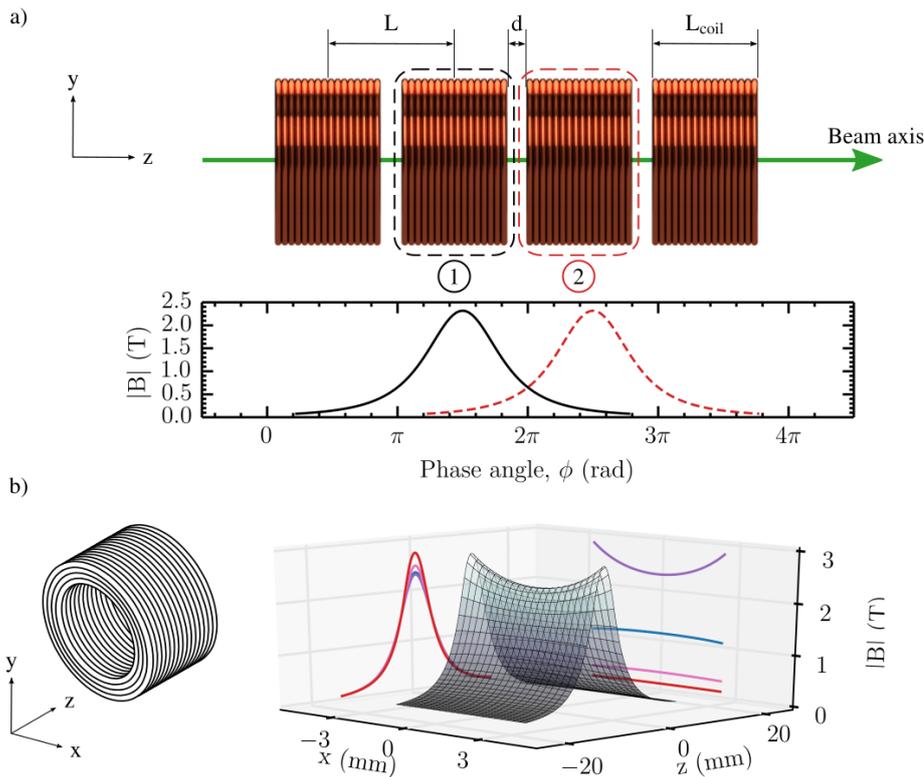


Figure 2.6: Schematic of a ‘conventional’ Zeeman decelerator. Two coils are shown, with the spatial periodicity defined in terms of the phase angle  $\phi$ . A plot of the magnetic field intensity in a plane through the centre of the decelerator is shown; the field switches from focussing to de-focussing from the centre to the boundary between coils. Image reproduced from L. McArd [2].

In order to effectively operate a conventional Zeeman decelerator, it is necessary to determine a sequence of switching times for the solenoids which maximises the fraction of the beam that is successfully loaded into the machine and decelerated. The switching times are traditionally calculated in terms of an idealised ‘synchronous atom’ which travels along the decelerator axis, with the ‘phase angle’  $\phi$  being defined as its position relative to a pair of coils;  $0^\circ$  being equidistant between them, and  $90^\circ$  being at the centre of the active coil. The times are determined such that the synchronous atom reaches the same phase angle when each active coil shuts down, so that it loses the same amount of kinetic energy at each step. Particles (in LFS states) that do not have the same position and velocity as the synchronous atom will oscillate around it in phase-space, due to either being decelerated more or less than it at each step. The volume of phase-space which particles can occupy and still remain in an orbit around the synchronous atom are defined as the spatial and velocity acceptance, which are used in characterising decelerators. Maximising the acceptance and the reduction of the kinetic energy of the beam requires finding the optimal value of the phase angle, which in turn depends on the beam species, flow velocity and translational temperature. Determining these timing sequences can be difficult; a modern approach is to simulate the system with evolutionary algorithms, as demonstrated by Toscano *et al.* [167].

A problem with conventional Zeeman decelerators is that the transverse field switches from focussing to defocussing (concave to convex) and back again between coils; this couples the longitudinal and transverse motions of the particles. Some groups have addressed this issue by adding additional transverse focussing, either by reconfiguring some of the decelerator coils [49] or adding hexapole magnet sections [50]. Other solutions attempt to provide roughly constant longitudinal and transverse confinement to effectively produce a moving three-dimensional trap, an approach successfully applied to Stark decelerators [36]. Examples of such ‘moving-trap’ or ‘travelling-wave’ decelerators include the designs described in Trimeche *et al.*, utilising flattened helical coils [51], Lavert-Ofir *et al.* [52], composed of many pairs of cylindrical solenoids in anti-Helmholtz configuration, or Damjanović *et al.* [53, 168], with a novel double-helical geometry. Ji *et al.* have recently proposed a machine based on the Lavert-Ofir design but with different current control [55]. Our travelling-wave Zeeman decelerator utilises coils based on those in the Trimeche machine [1, 2]. Comparisons of these decelerators to our experiment are given in sections 2.3.3 and 4.6.

### 2.3.3 Travelling-Wave Zeeman Decelerators

There are several designs of ‘moving-trap’ or ‘travelling-wave’ Zeeman decelerator. In this section their key features will be briefly described, and compared with the experiment constructed in Durham by the Carty group. Some comparison of experimental and theoretical performance will be given in section 4.6. The Durham decelerator was based on the machine described by Trimeche *et al.* [51] (henceforth referred to as the Paris decelerator), which utilises flattened helical deceleration coils. Our coils are described in section 2.3.4 (and more completely in McArd [2]). These offer the advantage of strong longitudinal trapping fields, and minimise any asymmetry arising from the wiring connectors, which can be kept away from the beam tube. They also give good access along the  $x$ -axis which can be useful if additional vacuum pumping is needed. The disadvantage of this design is the weak confinement in the transverse directions, especially in  $x$ , necessitating an additional focussing system to prevent radial particle losses.

The Durham coils have the same basic layout as the Paris coils, but have several key changes. The Paris design had crossing angle  $\alpha = 72.1^\circ$ , wire diameter  $d_w = 354\mu\text{m}$  and spatial periodicity  $\lambda = 11.6\text{mm}$ . There were two coil modules, each comprising 12 periods and 139.2mm in length. The coils were made of 16 wires connected in two 8-wire phases. As described in section 2.3.4 the Durham design has  $\alpha = 70^\circ$ ,  $d_w = 1.2\text{mm}$  and  $\lambda = 40.9\text{mm}$ . Our coils are composed of 3 periods for a module length of 122.6 mm and there are 4 modules. Our coils are also formed of 16 wires but these are grouped into 4 phases. This reduces changes in the shape of the trap field as it propagates. The coils are also paired above and below the sample tube to further increase the field strength. The crossing angle was adjusted to maximise the longitudinal trap depth and increase the spatial extent of the trap.

The sinusoidal currents applied to the deceleration coils in the Paris decelerator are produced by an H-bridge circuit with pulse-width modulation, which was also adopted for the Durham machine (see sections 2.3.5 and 2.3.6). The Paris power supplies can generate a peak current of 550 A amplitude whereas the Durham supplies have been designed to handle up to 1000 A, which is possible due to the larger gauge wire we have used. The electrical properties of coils with different numbers of periods and wire diameters were assessed; it is desirable to make the individual wire sections as long as possible to minimise the number of wire connections necessary, but if they are too long then they can have a large voltage drop across them. For details of the design process see Mizouri [1] and McArd [2]. Trimeche *et al.* report generating a magnetic wave moving at  $464\text{m s}^{-1}$  but to achieve this velocity the peak current is only 250 A. This is in line with our findings as reported in section 4.2; for example, at  $350\text{m s}^{-1}$  we find that the maximum current amplitude we can achieve is  $\approx 400\text{A}$ . As noted, however, our paired coils give a somewhat stronger field for a given current than the Paris design.

The evacuated tube through which the molecular beam passes in the Paris experiment has a relatively small diameter of 1.2 mm, which constrains the spatial acceptance. To improve this we utilised a tube of 5.2 mm diameter. The Paris experiment used a General valve with electric discharge excitation, we instead have an Even-Lavie valve with dielectric barrier discharge (DBD, see sections 2.2.2 and 2.2.3). The Even-Lavie valve can produce shorter pulses than the General valve, of order of  $\approx 20\text{--}30\ \mu\text{s}$  duration compared to  $\approx 100\ \mu\text{s}$ . The DBD produces more efficient excitation with less heating than the discharge fitted to the General valve, however the latter was reported as being able to produce short excitation pulses of  $2\ \mu\text{s}$  duration [51]. This means that the population of excited atoms loaded into the decelerator can still have a relatively small temporal width.

The Paris system also utilised a wire quadrupole to add an additional lateral focussing field. This was composed of 16 straight wire segments, 4 on each corner. The Durham decelerator was constructed with a very similar quadrupole, described in section 2.3.7. Our quadrupole only had 2 segments per corner but was designed to operate at higher currents, of  $\approx 700\ \text{A}$  compared to  $\approx 130\ \text{A}$ . The wire quadrupole and its proposed permanent magnet replacement are discussed in more detail in chapter 4.

Another moving-trap decelerator design was proposed by Narevicius *et al.* [169] and described in Lavert-Ofir *et al.* [52] (henceforth referred to as the Tel Aviv decelerator). This comprises a series of cylindrical solenoids, connected in overlapping anti-Helmholtz pairs. Initially, there were 213 traps in a 1.1 m deceleration stage, this was later expanded to 480 traps over 2.4 m [56]. The solenoids are constructed of 4 turns of 0.4 mm diameter wire, and the distance between neighbouring traps is 5 mm. The cylindrical geometry of the coils has the advantage of producing uniform trapping fields in all radial directions, unlike the flattened coils of the Paris and Durham decelerators. They have an inner radius of 10.2 mm, which fits closely around the large 9.5 mm beam tube. The Paris design was chosen as the basis for the Durham decelerator as it was believed that the coil geometry would introduce fewer oscillations in trap depth and position than the Tel Aviv geometry [1], however there does not seem to be a particular advantage in this regard. Trap oscillations are discussed further in section 4.2. The Tel Aviv decelerator uses RLC circuits to provide tunable current pulses with multichannel outputs to the deceleration coils [170].

A somewhat different approach is described in Damjanović *et al.* [53, 168] (henceforth referred to as the Basel decelerator). They report on the construction of a machine comprising two layers of 0.36 mm square wires wrapped in spirals around a cylindrical section; both the inner and outer layers are composed of 16 wires each and are wrapped with opposite handedness. A sinusoidal current is applied to each wire; the matching pair of wires in the inner and outer coils have the same phase, moving around the cylinder each subsequent pair have a temporal phase offset of  $\pi/8$ . The wire spirals have a periodicity of  $\lambda = 14\ \text{mm}$ , and the tube through which the molecular beam passes has a diameter of 4 mm. There are 16 coil modules, each of which is 56 mm in length, so that the deceleration stage is 896 mm long in total (the authors state that they plan to add a further 16 modules). The Basel decelerator also differs in having the complete coil assemblies inside vacuum. The AC current supplied to the coils is synthesised from a DC supply via H-bridge circuits in a similar way to the Durham decelerator, however as well as pulse-width modulation the technique of modified square-wave is employed [168]. The geometry of the traps produced is rather similar to those of the Paris or Durham decelerators, namely strong confinement in the longitudinal dimension and one of

the transverse directions, but relatively open along the other transverse axis; for example compare figure 2 in Damjanović [168] with the figures in section 2.3.7. In the Basel system, however, the traps effectively rotate around the beam axis as they propagate. With suitable experimental parameters this can mitigate losses along the weak axis.

The Tel Aviv decelerator has a peak current of 500 A, at which the forward trap depth is reported as being  $\approx 0.8$  T. For comparison, in the data presented in chapter 3, the Durham decelerator was operated at 400 A which gave rise to longitudinal trap depths of  $\approx 0.5$  T. When operating at its design maximum of 1000 A the longitudinal depth would be  $\approx 1$  T, although the peak current that can be achieved is dependent on the wave velocity. (As discussed in chapter 5 it is not always desirable to operate at the maximum current. This is because doing so results in a broader range of velocities in the decelerated packet, which can then be harder to trap post-deceleration). The Paris decelerator was reported as being operated at a peak current of 250 A when guiding metastable  $^3\text{P}_2$  argon [51]. The trap depth was given as a maximum velocity of  $15.5 / 11 \text{ m s}^{-1}$  for atoms in the  $m_J = 2 / 1$  states, which corresponds to a field strength of  $\approx 0.3$  T. (At 250 A the Durham decelerator trap depth is  $\approx 10\%$  greater than this). A typical peak current for the Basel decelerator when operating with trap velocities of  $\approx 450 \text{ m s}^{-1}$  is 300 A, with a corresponding trap depth of  $\approx 0.4$  T [168].

Assessing the capabilities of the different decelerators with regard to each other requires more than simply comparing the spatial extent and depth of the moving traps they produce. A common metric for the performance of a decelerator is its phase-space acceptance, the six-dimensional volume in position and velocity space occupied by particles which remain stably clustered within a moving trap. This is often estimated via the use of Monte Carlo methods, which will be discussed in chapters 3 and 4. Some comparisons between the Durham decelerator and the other experiments mentioned in this section will be given in section 4.6. The remainder of this chapter will give further details of the components of the Durham decelerator.

### 2.3.4 Decelerator Coils

Our travelling-wave Zeeman decelerator uses flattened helical coils which will be briefly reviewed here, for more detail on the design and construction of our coils see McArd [2]. The coils are based on those of the Paris decelerator [51]; technical differences between the Durham and Paris decelerators is given in section 2.3.3, with comparison of some experimental results in section 4.6. Each module of the decelerator is composed of two pairs of coils, one above and one below the sample delivery tube; the upper and lower coils are wound with opposite handedness. This geometry offers the advantage of physical access along the  $x$ -axis, which can be useful in long decelerators if additional vacuum pumping is necessary, and also produces stronger trapping fields for equivalent current than conventional solenoids.

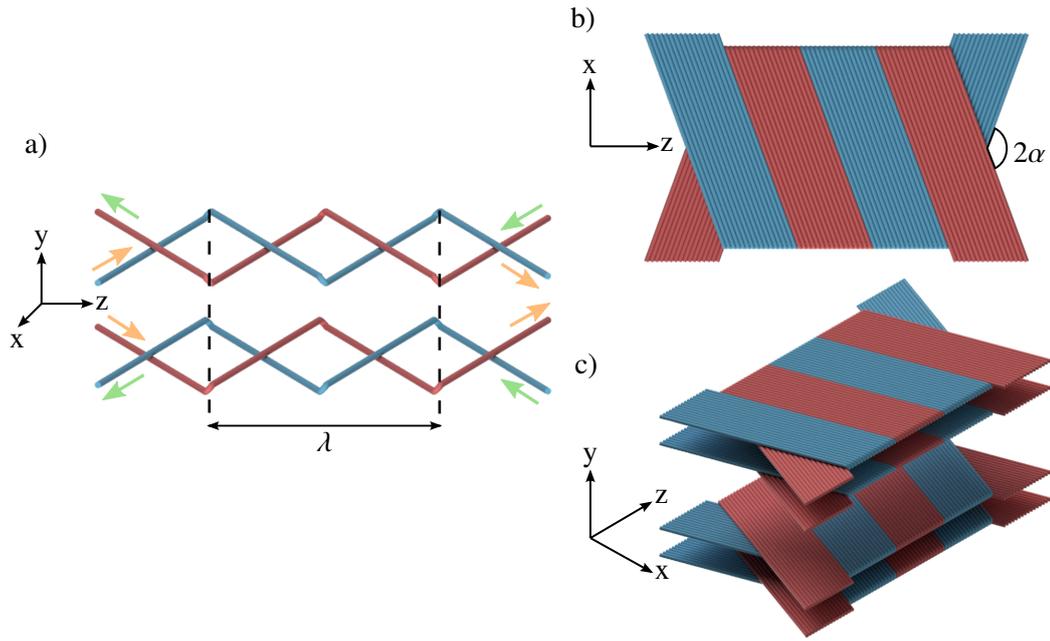


Figure 2.7: Structure of the helical coils used in the Carty group travelling-wave Zeeman decelerator, based on the design of Trimeche *et al.* [51]. a), a simplified view of two coils,  $\lambda$  is the spatial period. b), a complete coil comprising two sets of 16 wires.  $\alpha$  is the half angle at which they cross. c), a complete module containing two pairs of coils wound with opposite handedness. The opposing current flows generate a periodic anti-Helmoltz array. Image reproduced from McArd [2].

The operation of the coils can be easily understood by looking at the simplified schematic in figure 2.7a), which depicts a pair of wires from each of an upper and lower coil, folded to make an angle  $\alpha$  with the  $z$ -axis. The current flows in opposite directions in the red and blue wires, so that each coil can be viewed as a series of individual rhomboid current loops. The upper and lower loops also rotate in opposite directions, so that the magnetic field produced is analogous to that of three pairs of anti-Helmholtz coils lined up along the  $z$ -axis. Figure 2.7b) shows the complete coil; both the red and blue circuits contain 16 wires. Figure 2.7c) shows a complete module, with each coil doubled up to increase the trapping field. The sample delivery tube runs between the pairs, along the  $z$ -axis. The spatial periodicity of the coils (and thus also the field) is  $\lambda$ , given by:

$$\lambda = \frac{32d_w}{\sin(\alpha)}, \quad (2.3.10)$$

where  $d_w$  is the wire diameter. The angle  $\alpha = 70^\circ$  and the wire (kapton coated AWG 18 copper) has diameter  $d_w = 1.2$  mm, so that  $\lambda = 40.9$  mm. For clarity the coils shown are two periods long, the actual units in the experiment are three periods long giving a total length of 12.3 cm. The magnetic traps are made to ‘move’ along  $z$  as required by applying the appropriate time-dependent current to the coils. Each of the blue and red wire bundles is separated into four phases of four wires each, and the time-varying currents are applied with suitable time offsets. This makes the trap propagation more smooth, at the expense of making the power electronics and control systems more complex; further details are given in sections 2.3.5 and 2.3.6.

Figures 2.8 and 2.9 show the calculated field magnitudes in the  $xz$  and  $yz$  planes for a constant velocity guiding sequence at  $335 \text{ m s}^{-1}$ , with a peak current of 400 A. Due to the coil geometry the traps are more strongly confining in the  $y$  direction than in  $x$ ; see section 2.3.7 for details on adding transverse focussing. For more on simulating the decelerator fields see section 3.6.1.

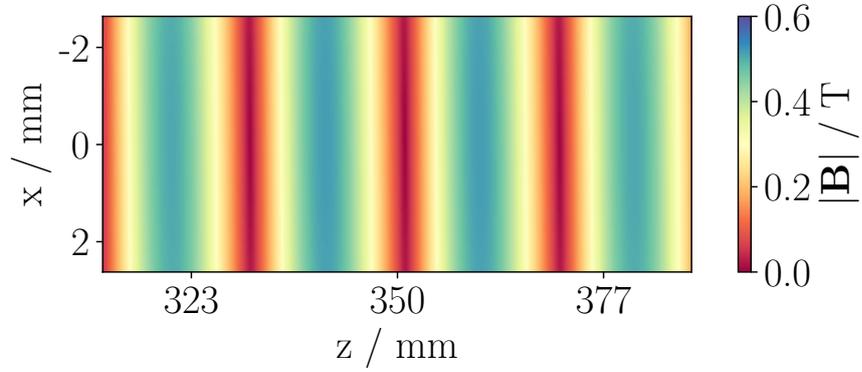


Figure 2.8: Calculated decelerator fields in the  $xz$ -plane across several traps, the bounds of the vertical axis are the internal diameter of the sample delivery tube. Coils are running with a peak current of 400 A. The traps are not very strongly confining in the  $x$ -direction.

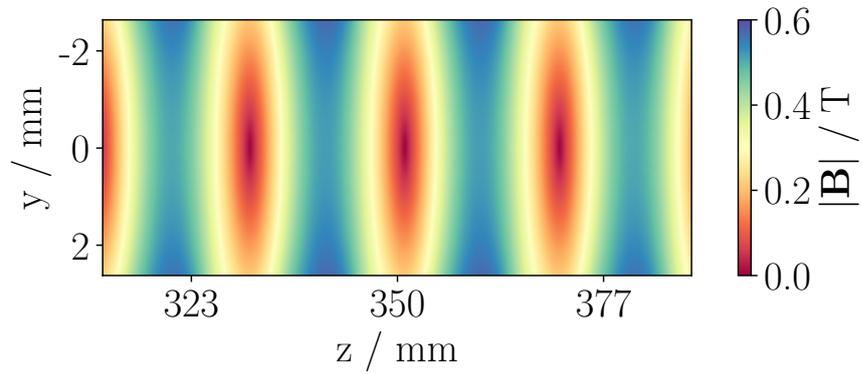


Figure 2.9: As per figure 2.8, but in the  $yz$ -plane. The decelerator traps are more strongly confining in  $y$  than in  $x$ .

### 2.3.5 Decelerator Electronics

This section will give only a brief outline of the power electronics used in the travelling-wave decelerator; for a detailed description see McArd [2]. Section 2.3.4 described the layout of the decelerator coils, each of which is made up of wires connected in four phases. To generate the required magnetic fields the individual phases must be supplied with (peak-to-peak) sinusoidal currents of up to 1000 A with a frequency of up to 10 kHz, with a phase shift of  $\pi/4$  radians between them. AC power supplies are problematic for this application because of the significant heat generated in resistive devices at such large currents. Instead, we utilise DC supplies in conjunction with fast switching to generate the sinusoidal current waveforms via pulse-width modulation (PWM, see section 2.3.6). The power electronics were constructed in-house, with the switching based around ‘H-bridges’, a circuit design which is used to reverse the polarity of a voltage applied to a load. The H-bridges are built around insulated-gate bipolar transistors (IGBTs); commonly found in applications such as welding equipment, these are efficient, fast and capable of handling large currents. These allow the entire supply voltage to be rapidly applied across a coil phase in either direction. Each coil module is connected to four H-bridges, one per phase, which connect to a bank of three pairs of 2.2 mF capacitors, fed by a DC power supply capable of delivering 800 V and 500 A.

### 2.3.6 Pulse-Width Modulation

The decelerator coils can be treated as simple RL-circuits comprised of an ideal resistor in series with an ideal inductor. When a source of emf  $\mathcal{E}$  is applied the current  $I$  through the coil rises exponentially with a characteristic time  $\tau = L/R$ , where  $L$  and  $R$  are the inductance and resistance, asymptotically approaching a maximum  $I = \mathcal{E}/R$ . When the applied voltage is removed, the current flow decays, also exponentially. The rise and fall can be approximated by the equations:

$$I(t) = \frac{\mathcal{E}}{R} (1 - \exp(-t/\tau)), \quad (2.3.11)$$

$$I(t) = \frac{\mathcal{E}}{R} \exp(-t/\tau). \quad (2.3.12)$$

These expressions allow us to calculate the voltage pulses needed to produce the desired current waveform, via pulse-width modulation (PWM). PWM is a technique for varying the average voltage and current fed to a load by rapidly switching the supply on and off. PWM is suited to ‘inertial’ loads which react more slowly to changes; this applies to solenoids because of Lenz’s law. Application to the decelerator is shown in figure 2.10. For a given velocity the frequency of the wave is determined from the spatial period of the coil (*cf.* section 2.3.4). Combined with the desired peak current this defines the reference waveform shown in blue. A tolerance envelope is determined by choosing a maximum value by which the achieved peak current can vary from the target, typically of order  $\approx 5 - 10\%$ . In an idealised case this tolerance could be set to an arbitrarily small value and the output waveform would approach the reference one arbitrarily closely, but this is constrained by the rise and fall times of the current. In general, the faster the wave velocity, the larger the tolerance value must be.

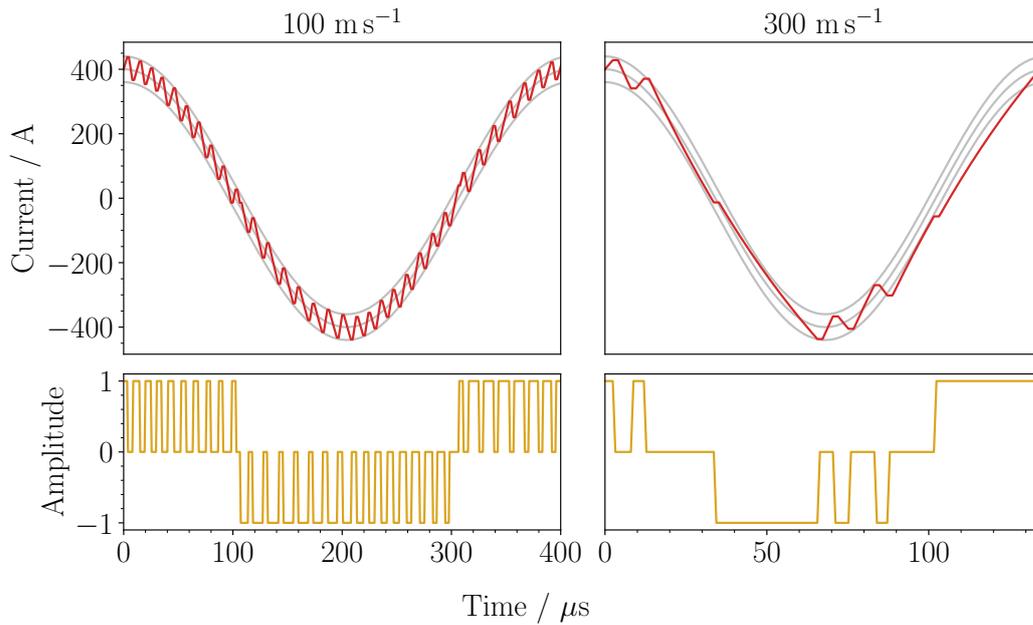


Figure 2.10: Example sinusoidal current sequence generation for travelling-wave velocities of (left) 100 and (right) 300  $\text{m s}^{-1}$ , decelerator running in guiding mode with a peak current of 400 A. Top row; a reference waveform is calculated with a tolerance envelope defined as a fraction of the peak current (grey), in this case 10 %. The effective synthesised waveforms generated by the applied voltage pulses across the coil are in red. Bottom row; the applied voltage pulses across the coil.

The pulse sequences are calculated as follows; the voltage is applied until the current through the coil reaches the upper (*i.e.*, largest magnitude) limit of the tolerance envelope. The voltage is removed and the current drops until it reaches the lower limit. This process is repeated as necessary, and is shown in the upper row of figure 2.10; once the waveform crosses the origin the polarity of the voltage is reversed. The entire sinusoidal waveform is built up using the fixed-amplitude variable-width pulses shown in the lower row. The two examples shown are for fixed-velocity guiding, for the case of deceleration the same approach is used, but the frequency of the reference waveform is linearly reduced. The rate of this ‘chirping’ down is chosen so that the first full trap to form is at the target velocity as it reaches the end of the last coil. Additionally, when calculating the pulse sequences the finite recovery time of the switches and any response or delay times of other electrical components must be considered and factored in to the calculations. The viability of a particular waveform to be encoded with PWM is determined by the wave frequency and amplitude (because of the current rise and fall times). If a viable solution is found, the pulse sequence is programmed into an FPGA which connects to the power electronics via fibre-optic couplers. An example of the expected current waveforms determined for each phase of the four-module decelerator is given in figure 2.11 - this is a simulation, but the achieved current has been experimentally measured with a clamp meter, and found to agree reasonably well with predictions, see figures 4.14 and 4.15 in section 4.2. For details on the development of these systems see McArd [2].

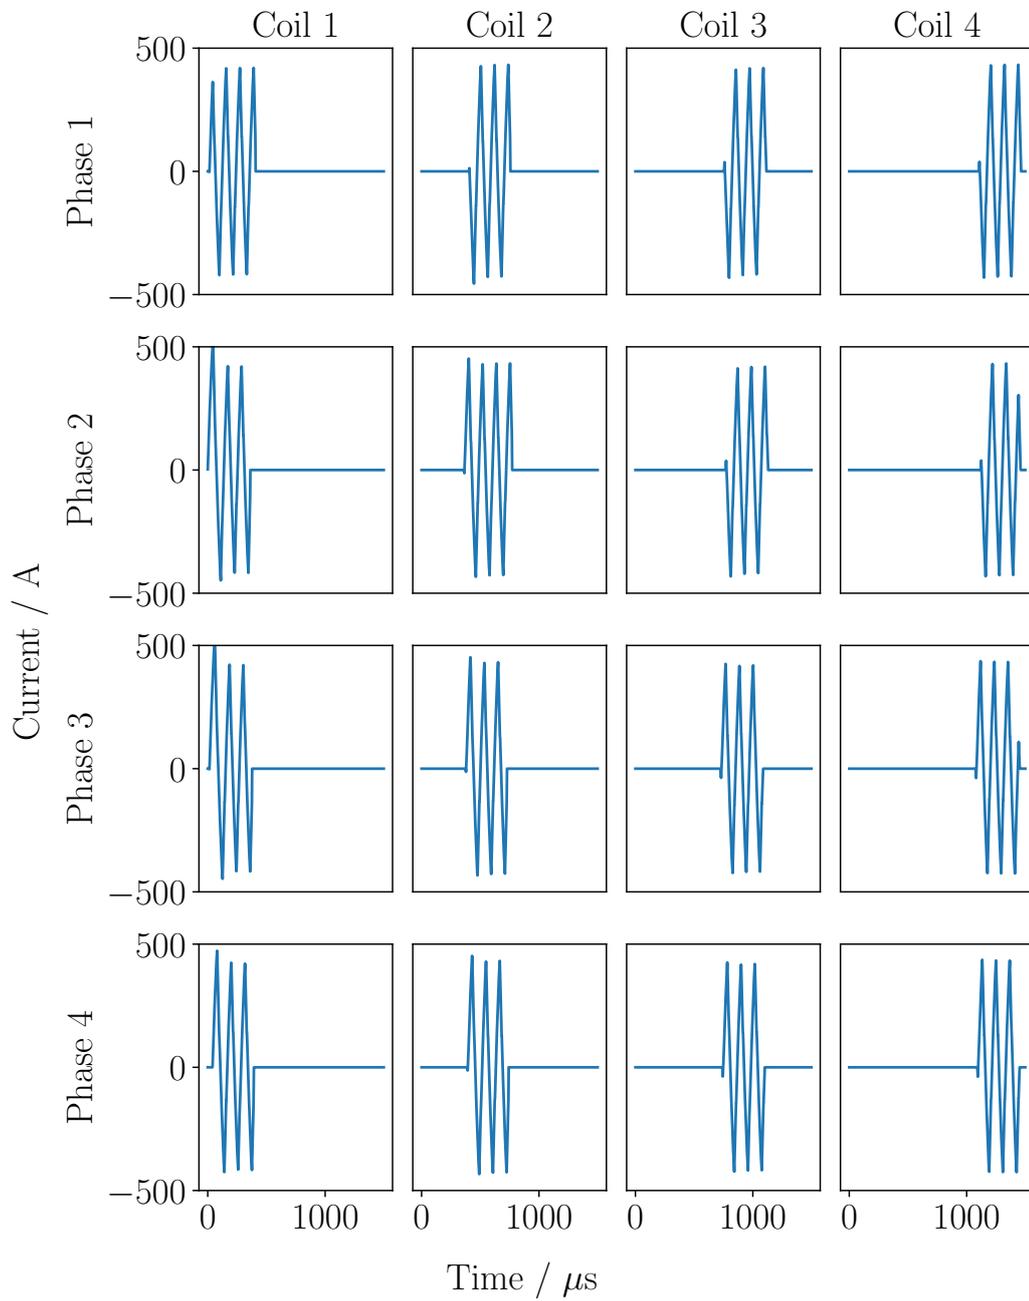


Figure 2.11: The simulated current through each phase of the four-module decelerator as a function of time, for constant-velocity guiding at  $350 \text{ m s}^{-1}$ . Target peak current is 500 A driven by 650 V.

### 2.3.7 Quadrupole

As discussed in section 2.3.4, the decelerator fields are not very confining in the transverse directions, especially in the  $x$ -direction. To improve this, an additional quadrupole field was added to the decelerator. This was originally provided by a low-voltage, high-current DC circuit running in two wires per corner, as shown in cross-section in figure 2.12. Typical operating parameters for this were  $\approx 40$  V and  $\approx 700$  A. This design provided a strong transverse focussing effect, but was problematic for several reasons. It generated significant heat and thus required active cooling, and was not easily extendable for a longer decelerator without adding additional high-output power supplies and fast-switching high-current contactors. Any failure of a contactor would result in the circuit latching closed and burning out almost immediately. Also, since the quadrupole was composed of one continuous circuit the wire connectors at the end of each corner were orthogonal to the beam path, and produced significant unwanted field components in the  $z$ -direction. Mitigating this necessitated very fast switch on and off times, further adding to the complexity of the system.

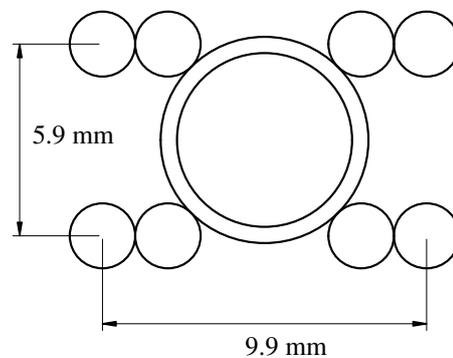


Figure 2.12: Wire quadrupole, the wires on each corner are paired. The current flow is in the same direction for diagonal corners. The sample delivery tube is in the centre. This design provides an effective transverse focussing field that is uniform along the length of the decelerator, although the power supplies and control gear are complex.

An alternative quadrupole was designed based on custom-made permanent magnets. These were shaped to fit into the slots in the decelerator cooling blocks that the wire quadrupole formerly occupied, as shown in figure 2.13. As well as simplifying the control and power requirements of the decelerator this design provides significantly stronger focussing fields. The magnets are made from N50 grade NdFeB with magnetisation vectors as shown in the figure, with two different directions oriented at  $30^\circ$  from the horizontal. For details of the wire quadrupole electronics and the design of the permanent magnets see McArd [2].

In order to compare the two solutions their field profiles were calculated. The wire quadrupole is simple to model along most of its length; similarly to the decelerator fields in section 3.6.1, the Biot-Savart law is used to treat each of the eight wire segments, and the resultant individual components can be summed together. The permanent magnet quadrupole has no analytic solution, and was originally modelled in Radia [171]. This has problems dealing with curved magnetic surfaces, so a simplified model of the transverse field was created using the FEMM software package [172], which made the assumption of uniformity along the beam axis. (See section 3.6.2 for Radia and FEMM). The assumption of uniformity along the beam axis is not strictly valid; where the individual magnets

adjoin (every 30 mm) there will be a small, periodic variation. This was shown to be very small using the Radia model. The transverse field magnitudes of the two systems are compared in figure 2.14, and it can be seen that the permanent quadrupole has a very similar profile with stronger focussing compared to the wire quadrupole.

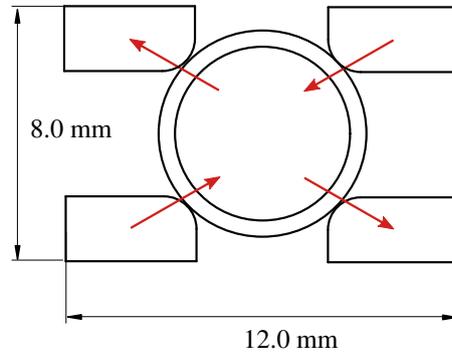


Figure 2.13: Permanent magnet quadrupole. Custom shaped pole pieces replace the paired wires seen in figure 2.12. The sections have two possible magnetisation vectors, shown by the red arrows.

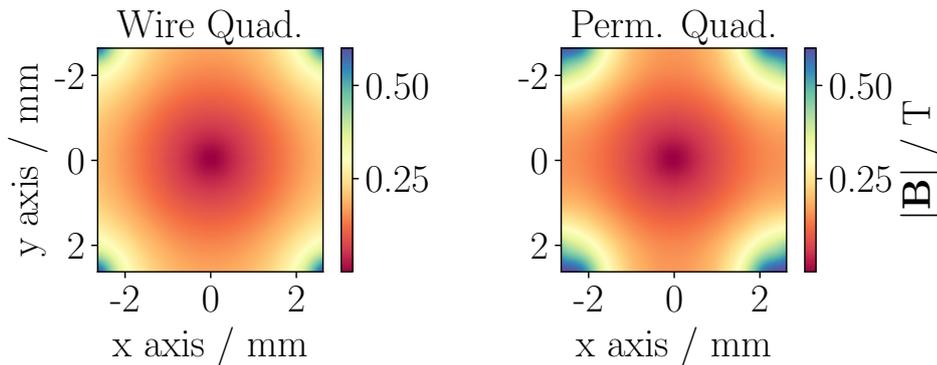


Figure 2.14: Comparison of the transverse focussing magnetic field magnitudes for the wire (left) and permanent magnet (right) quadrupoles. The permanent system provides a very similar field profile and stronger focussing compared to the wire, with reduced complexity and cost. However, the field has a small periodic variation along the length of the decelerator, due to the boundaries between pole pieces.

Although both designs are practically uniform along the length of the beam axis, they interact differently with the decelerator fields. The latter are shown as plots in the  $xz$ - and  $yz$ -planes in figures 2.8 and 2.9; figure 2.15 below shows the effects of adding the quadrupole field. Although it generally improves the trap confinement in  $x$ , which is where the decelerator field is weakest, the wire field worsens confinement in  $y$  for alternate traps.

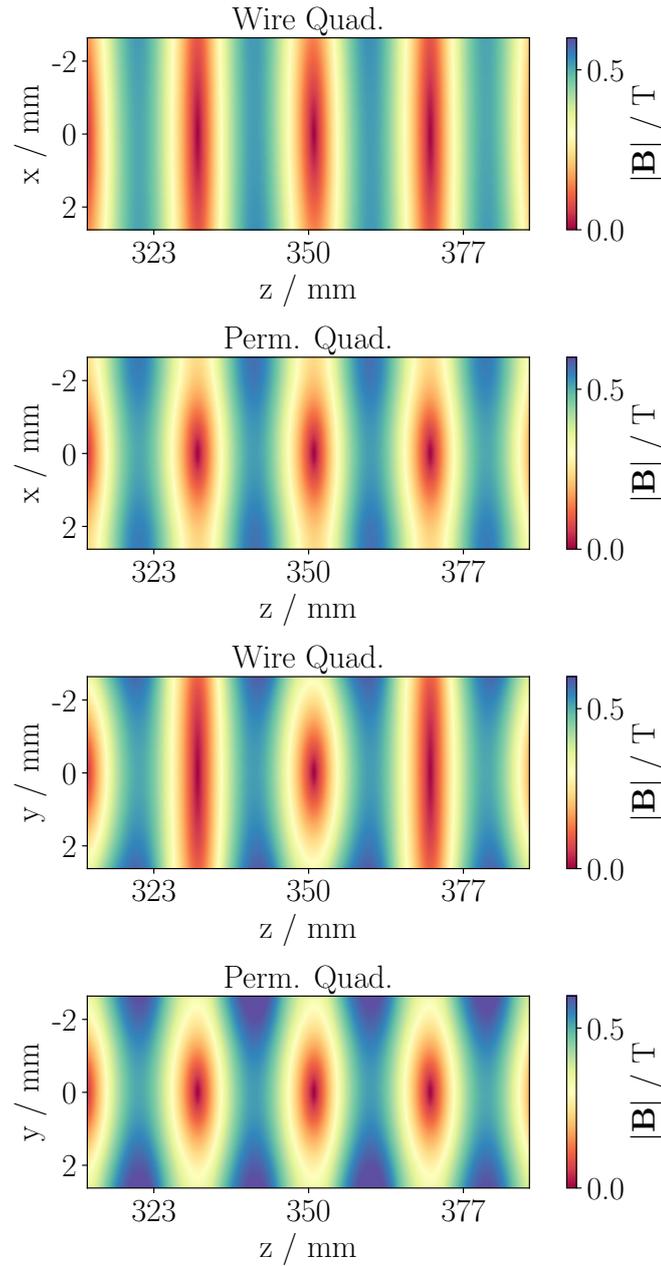


Figure 2.15: Comparison of decelerator field plus wire or permanent quadrupole, in the  $xz$ - and  $yz$ -planes. Compared to the decelerator field alone in figures 2.8 and 2.9 it can be seen that the wire quadrupole improves or worsens alternate traps, due to the relative geometry of the fields. The permanent magnet field does not have this problem and provides a much better transverse focussing effect.

This effect arises due to the orientation of the decelerator and quadrupole fields. For the wire design, the flux lines enter and leave in the horizontal and vertical directions, whereas for the permanent magnet system they are along the diagonals. This can be seen in figures 2.16 and 2.17, which show plots through the  $xy$ -plane along with the field profile along the  $x$ - and  $y$ -axes. The two longitudinal distances shown,  $z = 330$  and  $351$  mm, roughly correspond to the first two complete traps in figure 2.15, which are defocussing / focussing for the wire case.

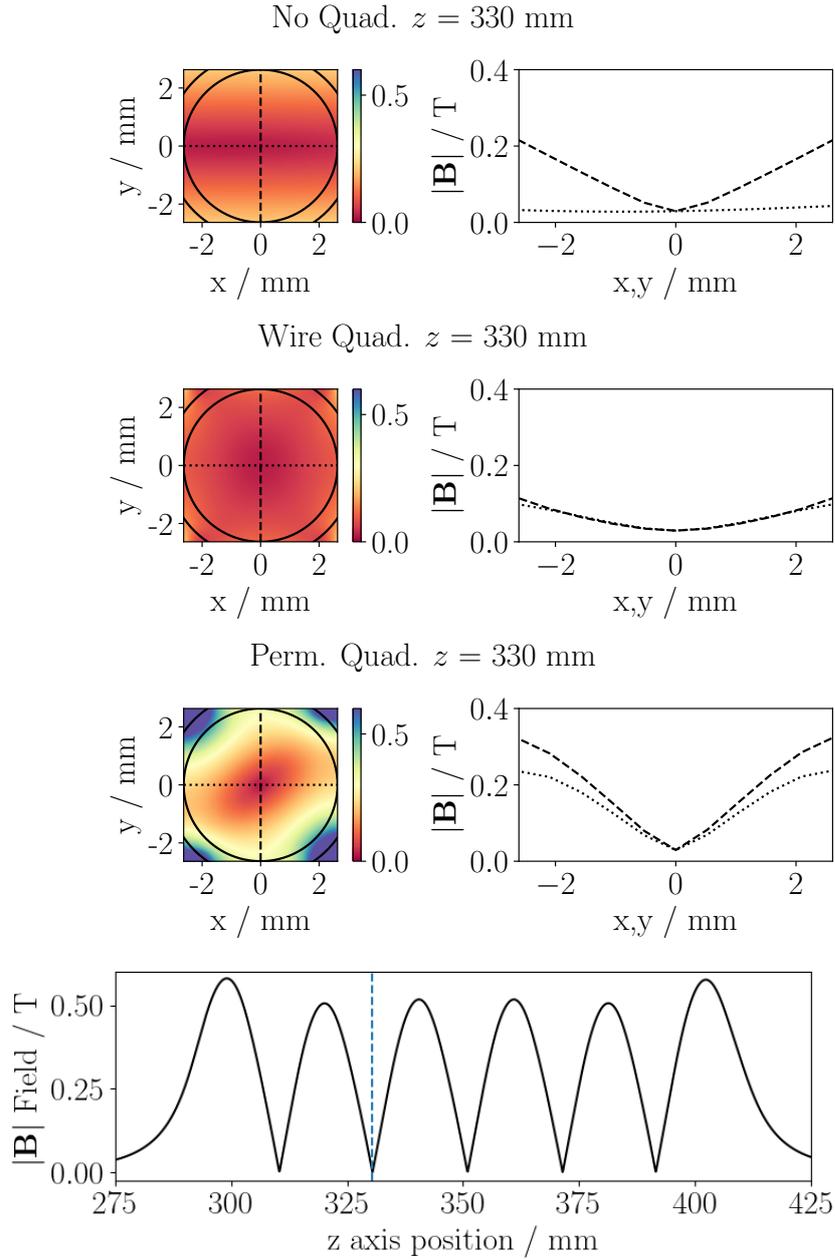


Figure 2.16: Field magnitude in the  $xy$ -plane at the centre of the trap at  $z = 330$  mm, bottom. As labelled, decelerator field alone, added wire quadrupole, added permanent magnet quadrupole. Left, solid circles show the inner and outer radii of the sample delivery tube. Left and right, the dotted / dashed lines show the  $x$ - and  $y$ -axes, the field profiles of which are shown on the right. Due to their relative geometry the wire quadrupole weakens the transverse decelerator field for alternating traps, as shown here. The permanent magnet field does not have this problem, because it has smaller components in  $x$  and  $y$ , being aligned along the diagonals, whilst still presenting a large field gradient.

Simulations of the entire decelerator were developed for both quadrupole designs, allowing them to be compared. Unfortunately, the data shown in this thesis were taken after the wire system was removed and before the permanent one was installed. Earlier data incorporating the wire system [2] are, however, in agreement with our predictions. The acceptance of the decelerator is discussed in more detail in chapter 4.

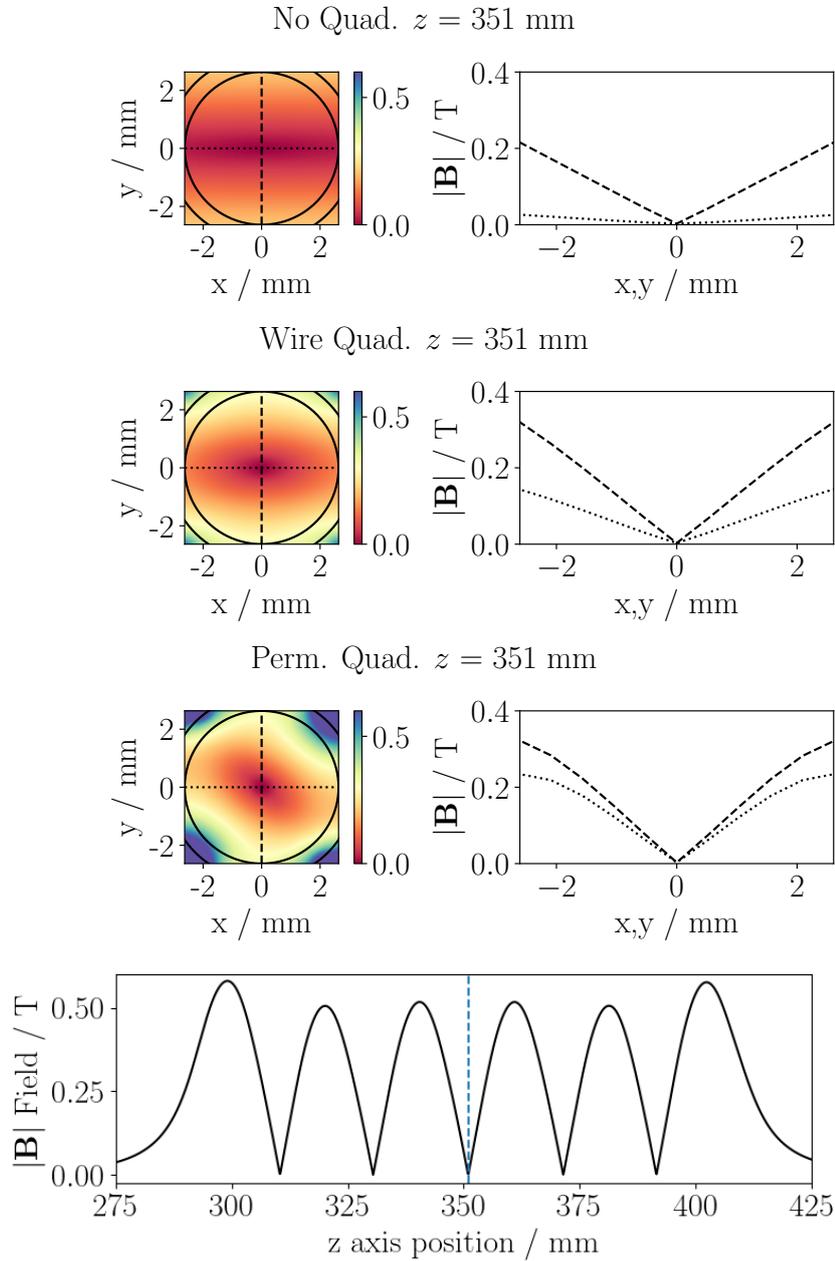


Figure 2.17: As per figure 2.16 but for the trap at  $z = 351$  mm. In this trap the wire quadrupole adds to the decelerator field to increase transverse focussing. The permanent magnet quadrupole is focussing for all traps, although it can be seen to be ‘handed’ in its interaction with the decelerator field.

## 2.4 Beam Detection

We have discussed the generation and propagation of a molecular beam; this section will cover the apparatus used to detect the beam, and theoretical models for the signal at the detector.

### 2.4.1 Time-of-Flight Profiles

Molecular beams may be characterised by an analysis of the signal they produce at a detector. Although various detection methods exist the basic principle of inferring the distribution of molecular velocities from the range of arrival times is common. We utilise an Even-Lavie pulsed valve source (see section 2.2.2) but continuous-expansion beams may also be subject to time-of-flight analysis, by introducing a chopper to selectively produce discrete packets of particles. In order to discuss the form of the expected signal, we begin with the case of an ideal gas in a container, and the flux of particles hitting a small region on the container wall. The fraction of particles with velocities within a volume element of velocity space  $d^3v$  centred on  $\mathbf{v}$  is given by the Maxwell-Boltzmann distribution:

$$f_S(\mathbf{v})d^3v = \left(\frac{1}{2\pi\sigma_S^2}\right)^{3/2} \exp\left(-\frac{v_x^2 + v_y^2 + v_z^2}{2\sigma_S^2}\right) d^3v, \quad (2.4.1)$$

Particles with speed  $v = (v_x^2 + v_y^2 + v_z^2)^{1/2}$  at an angle  $\theta$  to the wall can reach a circular region on the wall with area  $A$  in time  $\tau$  if they are located within a volume  $v \cos\theta A\tau$ . In spherical polar coordinates the volume element  $d^3v = v^2 \sin\theta$ , so for a number density  $n$  the flux  $\Phi_S(v)$  of particles reaching  $A$  in  $\tau$  as a function of speed is:

$$\Phi_S(v) = nA\tau \left(\frac{1}{2\pi\sigma_S^2}\right)^{3/2} v^3 \exp\left(-\frac{v^2}{2\sigma_S^2}\right) \cos\theta \sin\theta. \quad (2.4.2)$$

The range of the azimuthal angle  $\phi$  is  $0 < \phi < 2\pi$ , that of the polar angle  $\theta$  is  $0 < \theta < \pi/2$  since only half the particles are moving in the direction of the wall; integrating over these angles and dividing by  $\tau$  gives the flux  $\Phi(v)$  per unit time:

$$\Phi_S(v) = nA\pi \left(\frac{1}{2\pi\sigma_S^2}\right)^{3/2} v^3 \exp\left(-\frac{v^2}{2\sigma_S^2}\right). \quad (2.4.3)$$

Instead of a patch on the wall, the target could be a small aperture through which the gas escapes as an effusive beam. In this case  $\Phi_S(v) = dI_0/dv$ , where  $I_0$  is the on-axis intensity of the beam, and  $\Phi_S(v)$  is approximated by:

$$\Phi_S(v) = Bv^3 \exp\left(-\frac{v^2}{2\sigma_S^2}\right), \quad (2.4.4)$$

where  $B$  is a constant. Section 2.2.1 compared the distribution of atom velocities in a stagnation chamber with those in a supersonic expansion; similarly, equation 2.4.4 can be used to derive equation 2.4.5, the flux of particles in a supersonic expansion as a function of  $v$ :

$$\Phi_T(v) = Bv^3 \exp\left(-\frac{(v - v_T)^2}{2\sigma_T^2}\right). \quad (2.4.5)$$

As in section 2.2.1, the subscripts S, T refer to the stagnation and terminal (supersonic) regimes. Compared to 2.4.4 the expansion in equation 2.4.5 has a translational temperature with a narrower

variance  $\sigma_T^2$ , *i.e.* it is ‘colder’, but is moving in the laboratory frame with a flow velocity  $v_T$ . In the case of a continuous beam, either expression can be integrated across all  $v$  to give the steady on-axis beam intensity  $I_0$  [173]. For a skimmed supersonic beam we assume that the off-axis intensity varies very little. A pulsed beam is necessary to use time-of-flight to determine the velocity distribution; this can be generated using a chopper or, as in our case, a pulsed valve (see section 2.2.2). Whatever the source, if it has temporal width  $\Delta_t$  then the number of atoms in a packet will be  $N_0 = I_0\Delta_t$ , with a distribution:

$$dN = N_0\Phi_T(v)dv. \quad (2.4.6)$$

The intensity as a function of speed  $I(v) = dN/dt$ , so:

$$I(v) = N_0\Phi_T\frac{v}{t}. \quad (2.4.7)$$

We assume that the time-dependent signal  $S(t)$  at the detector has a linear relationship to the number of particles arriving there, given by  $S(t) \propto I(v)/v$ , so that  $S(t)$  takes the form:

$$S(t) \propto \Phi_T(v)/t. \quad (2.4.8)$$

Equality 2.4.8 leads to an equation for  $S(t)$ , where  $\Delta_{SD}$  is the distance from the source to the detector and  $v_T$  is the beam flow velocity; this can be fitted to real data for the parameters  $B$ ,  $v_T$  and  $\sigma_T$ , and from the latter the temperature  $T_T$  can be calculated since  $\sigma_T^2 = k_B T_T/m$ :

$$S(t) = B\frac{\Delta_{SD}^3}{t^4} \exp\left(-\frac{(\Delta_{SD}/t - v_T)^2}{2\sigma_T^2}\right). \quad (2.4.9)$$

It was briefly alluded to earlier in this section that the beam source must have some temporal width  $\Delta_t$ , but the argument presented so far has made no allowance for any factors that may broaden or otherwise distort the shape of the signal that a molecular beam produces at the detector. Examples of such factors may include the profile of the slots in a chopper wheel, the finite opening and closing times of a pulsed valve (which may be asymmetric), the dimensions of a discharge excitation electrode, and the response times of electronics in a detector. In practice, for many experimental setups these terms are negligible [174], producing errors in the fitted velocity  $v_T$  and temperature  $T_T$  that are small compared to those arising from other uncertainties. If the overall length and time scales of an experiment are large compared to those of the source, then a model such as equation 2.4.9 may be adequate.

If an account must be made for extra factors, the standard approach is to extend the function for the expected signal by convolving it with additional functions for the broadening terms [175]. For example, if the particles appear in the beam at times  $0 \leq t_0 \leq \Delta_t$  with some distribution  $g(t_0)$ , then instead of the signal  $S(t)$  we get a signal of the form  $S'(t)$ :

$$S'(t) = \int_0^{\Delta_t} S(t - t_0)g(t_0)dt_0. \quad (2.4.10)$$

The new expression for the signal  $S'(t)$  is obtained by replacing the variable  $t$  with  $t - t_0$  and integrating the product  $S(t - t_0)g(t_0)$  across the range of  $t_0$ .  $S'(t)$  can then be convolved with additional terms as necessary.

### 2.4.2 Microchannel Plate Detector

The detection chamber of the Zeeman decelerator was designed to operate with different detection schemes as necessary, dependent on the species in the molecular beam. Details of the chamber design and construction are given in Mizouri [1] and McArd [2]. When working with metastable argon a microchannel-plate detector (Hamamatsu compact MCP assembly F12334-11) was used.

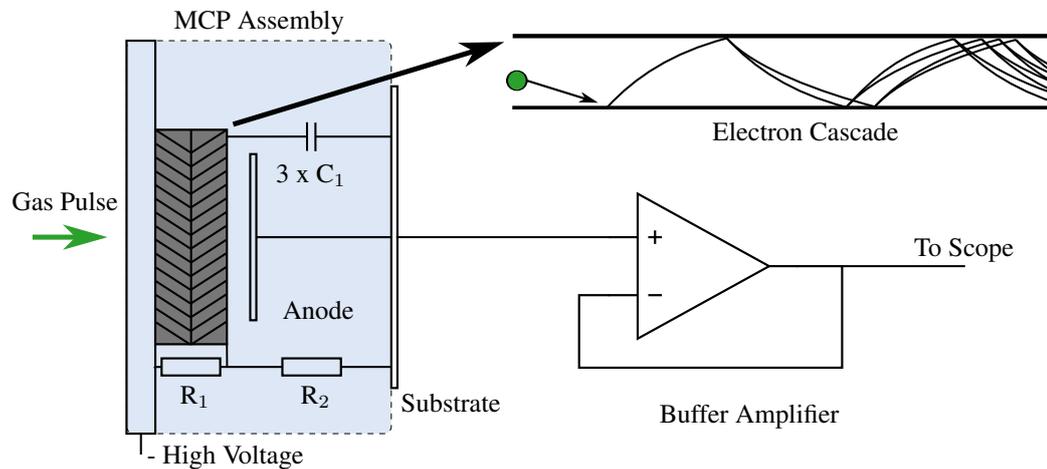


Figure 2.18: Microchannel-plate detector layout. Beam propagation direction is from left to right, shown are the two-section MCP with slanted channels, collection anode, high-voltage connection, signal output via buffer amplifier and an illustrative schematic of electron cascade within a channel. Image adapted from L. McArd [2].

The MCP comprises two plates containing millions of  $12\ \mu\text{m}$  diameter channels. High voltage is applied to the plates with a maximum magnitude of  $2.1\ \text{kV}$  but typically set to  $\approx 1.8\ \text{kV}$ . When a sufficiently energetic particle strikes the wall of a channel it liberates an electron, which is accelerated strongly by the electric field. This electron in turn collides with the channel wall producing a cascade, analogous to the operation of a photomultiplier. The electrons are collected at an anode, which is electrically connected to the plates via  $3 \times 150\ \text{pF}$  capacitors and a voltage divider where  $R_1$  and  $R_2$  are  $7.5\ \text{M}\Omega$  and  $0.2\ \text{M}\Omega$  resistors respectively, shown in figure 2.18. The MCP assembly is held within a spherical octagon on kapton-insulated groove-grabbers; if sufficient care is not taken to electrically isolate the unit then it can short out through the vacuum chamber. Due to the limited space within the octagon nuts and bolts could not be used, so the four mounting holes on the MCP body were tapped with  $\text{M}3.5 \times 0.6\ \text{mm}$  threads to accept nylon screws. The output of the anode is an SMA connector which is linked to a BNC electrical feedthrough on the chamber. The output impedance of the MCP is in the range  $20 - 80\ \text{M}\Omega$  which necessitates the use of a buffer amplifier to connect the signal output to the oscilloscope, see section 2.4.3. At room temperature and pressure of order  $10^4\ \text{Pa}$  the MCP has a typical full-width half-maximum response time of approximately  $1.5\ \text{ns}$ . Ground-state argon is not energetic enough to trigger the MCP, but the  $^3\text{P}_2$  state we aim to produce in the DBD (*cf.* section 2.2.3) has an energy of  $11.54\ \text{eV}$  above this [176]; see section 3.5 for further discussion of argon energy levels in our source.

### 2.4.3 Buffer Amplifier

The micro-channel plate detector (MCP, section 2.4.2) has output impedance between 20 to 80  $\Omega$ ; when connected to the oscilloscope (LeCroy WaveRunner 610Zi) which operates at 50  $\Omega$  or 1 M $\Omega$  this mismatch causes distortion to the signal. This is resolved by interposing a buffer amplifier with a 50  $\Omega$  output. The first version that was developed [1] was a unity-gain unit incorporating an OPA633 operational amplifier (op-amp) with the output connected to the inverting circuit; compared to the input, the output has a matched voltage but boosted current. This also reduces the impact of radio-frequency noise from the decelerator coils, which is only partially mitigated by the use of shielded cables. This design was reasonably successful but prone to fail due to overloading. As well as metastable atoms, the dielectric-barrier discharge (DBD, section 2.2.3) produces unwanted ions and ultraviolet radiation in a high-intensity burst which exceeds the safe operating voltage of the op-amp. This ‘ion peak’ can be seen in figure 2.21, from about 3  $\mu$ s to 200  $\mu$ s after the DBD trigger.

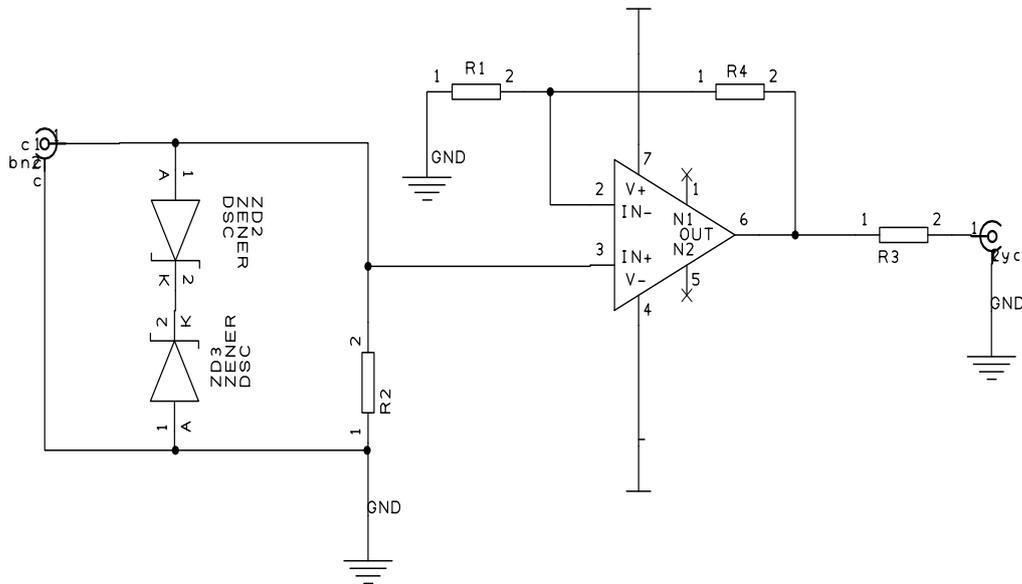


Figure 2.19: Schematic for the impedance-matching buffer amplifier used to connect MCP to oscilloscope. This was constructed in two versions, operating at 1 M $\Omega$  and 50  $\Omega$ ; the op-amps and values of resistors differ between these cases. See main text for details.

A replacement buffer amplifier was built based on the design described in Lohead [177], see figure 2.19. The Lohead unit, however, utilised a specialist THS3201 op-amp with a very high slew rate of  $6700 \text{ V } \mu\text{s}^{-1}$ , which is unnecessary for our experiment; for reference, the OPA633 only has a slew rate of  $2500 \text{ V } \mu\text{s}^{-1}$ . The THS3201 was therefore replaced with a TL081. The circuit contains back-to-back clamping diodes which limit the input voltage to  $2.7 \text{ V}$  in order to protect the op-amp. The resistors R1 and R4 were set at  $50 \text{ k}\Omega$  and  $500 \text{ k}\Omega$  to give a gain of 11. The  $1 \text{ M}\Omega$  resistor R2 was installed to stabilise the signal.

The TL081-based system worked well when coupled to the oscilloscope at  $1 \text{ M}\Omega$  but had insufficient power to drive a  $50 \Omega$  connection, which better reduces interference. The TL081 was swapped for an OPA552, with R1 and R4 set at  $50 \text{ k}\Omega$  and  $240 \text{ k}\Omega$  to give a gain of 5.8. The resistor R2 was found to operate best at about  $200 \text{ k}\Omega$ ; this was implemented as a trimmer to allow fine-tuning. The OPA552 has a maximum supply voltage of  $\pm 30 \text{ V}$  but under these conditions it may dissipate up to  $11 \text{ W}$ . The manufacturer also states that the performance of the chip will vary at different supply voltages and that it is safest to operate on the lowest feasible value. The device was tested with a variable power supply and found to work well at  $\pm 18 \text{ V}$ . Referring back to slew rates, both the TL081 and OPA552 have relatively low performance compared to the OPA633 ( $13 \text{ V } \mu\text{s}^{-1}$  and  $24 \text{ V } \mu\text{s}^{-1}$  respectively), but this turns out to be unimportant for our experimental setup.

The overall design for the buffer amplifier circuit can be easily adapted to other experiments with a need for faster response times by replacing the op-amp and adjusting the feedback resistors. The buffer was tested by feeding varying signals from a function generator into it and comparing the output with the input, for a range of input frequencies and amplitudes. Examples are shown in figure 2.20, with the input signal set to an amplitude of just over  $2.7 \text{ V}$ . This shows the effect of the voltage limiter, with the output clipping at  $2.7 \text{ V} \times 5.8 \approx 15.6 \text{ V}$ . It also demonstrates that there is no discernible distortion of the output due to the amplifier.

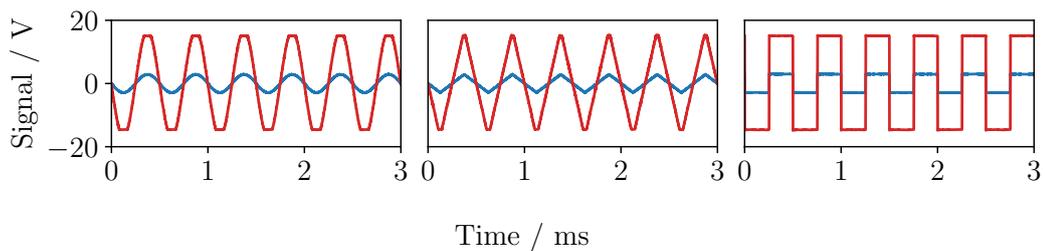


Figure 2.20: OPA552-based buffer amplifier tests with a function generator, blue trace, input signal, red trace, output signal. Input amplitude was chosen to be just greater than the maximum  $2.7 \text{ V}$  allowed by the safety clamping diodes, with a gain of 5.8 this means the output clips at  $\approx 15.6 \text{ V}$ . This can most clearly be seen in the sinusoidal trace on the left.

Figure 2.21 compares the detected signal at the MCP of argon cooled to  $\approx 139$  K for the old OPA633 (left) and new OPA552 (right) designs. The upper plots show a ‘transmission’ curve, in which the decelerator field is inactive, the lower both show constant-velocity guiding mode at 600 V / 300 A. It can be seen that the amount of ringing and noise on the signal due to interference from the operation of the decelerator is significantly reduced for the OPA552. Additional care was taken to provide a ‘clean’ earth and eliminate ground loops. The different temporal positions of the interference bands in the two plots are due to differences in the activation time of the decelerator electronics between the two experimental runs, which is done to adjust the loading of the molecular beam into the moving magnetic traps (see section 4.1.1).

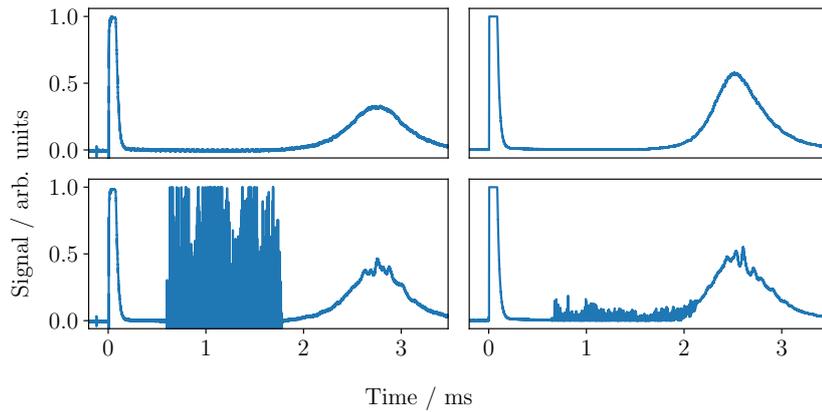


Figure 2.21: Comparison of buffer amplifiers. Top left: zero field, OPA633 unit. Bottom left: guiding, OPA633 unit. Top right: zero field, OPA552 unit. Bottom right: guiding, OPA552 unit. Both runs with valve at 139 K, decelerator at 600 V / 300 A and guiding at  $325 / 345 \text{ m s}^{-1}$  for left / right. It can be seen that the OPA552 amplifier more effectively reduces the interference produced by the decelerator power electronics and coils. The effect of the clamping diodes in limiting the voltage reaching the op-amp is visible in the clipped profile of the ion peak.

## 2.5 Experimental Parameters

This section describes how operating parameters for the experiment are chosen, specifically referring to the values used for the data presented in this thesis. It will be seen in chapters 3 and 4, however, that in order to analyse our experimental results and choose optimal parameters for the various components of the experiment it was necessary to develop a set of detailed computer models.

We begin with the beam source, the Even-Lavie valve with dielectric-barrier discharge, described in sections 2.2.2 and 2.2.3. In McArd [2] the valve was pulsed at 3 Hz; for the data in this report 2 Hz was chosen, to minimise the heat load arising from the excitation discharge. An additional source of heat, not present in the earlier work, was the electron-source filament introduced to improve the shot-to-shot stability of the discharge excitation (see section 2.2.4); this was active for all data described in this thesis. A current of 1.2 A was chosen, offering a good improvement in stability whilst being somewhat less than the rated maximum power dissipation of the filament. After several sets of data are taken, the filament is switched off to allow the valve to cool down (a pause of several minutes is in any case necessary in order to reprogram the decelerator electronics with new settings). Working with argon, a backing pressure of 6 bar was used. The system can work with much higher pressures [157, 158], but using the lowest pressure that gives an acceptable signal reduces wear on the gaskets in the Even-Lavie valve.

The parameters that can be adjusted in the source include the RF voltage, RF frequency, number of pulses, and valve opening time. The voltage can be set to between 1100 and 1200 V; 1150 V was chosen, again to compromise between signal and heat load. The frequency, number of pulses and valve opening all have to be chosen by trial and error; as the valve is cooled from room temperature all of these settings need to be manipulated. Due to differing ambient conditions, a working set of parameters from one day can not be relied upon to be good on the following day, but can serve as a baseline for adjustment. For the data presented in chapter 3 the valve and DBD were configured with a 25.6  $\mu\text{s}$  opening time, and 22 pulses at 650 kHz. The source is cooled with a target temperature of 140 K, but this will oscillate by several K during the experiment. Lastly for the source, the delay between the valve and DBD triggers is configurable. Again, this must be set by trial and error as the system cools and the flow velocity of the beam changes, and a value of 85  $\mu\text{s}$  was used once the source stabilised.

The decelerator parameters include the initial and (if decelerating) final wave velocities, the peak current and driving voltage, and the delay between the DBD trigger and coil activation. The maximum operating voltage of the deceleration coils is 800 V, but it is desirable to operate at a lower value in order to reduce wear on the electrical components. The peak current achievable is strongly related to the desired wave velocity, as discussed in section 2.3.6, with the wave velocity set according to the estimated flow velocity of the beam. Section 2.4.1 described the basic theory of fitting a time-of-flight profile to a detected beam, and this will be developed in more detail in chapter 3. When working with argon cooled to  $\approx 140$  K we will show that we find the beam to have a central velocity in the region of 350 to 400  $\text{m s}^{-1}$ ; to match this, our power electronics can realistically generate a peak current of  $\approx 400$  A. For the work in this thesis this was the chosen current, driven by 600 V - calculations show that using a higher voltage has only a small effect on the current in the coils. The choice of electronics delay setting is extremely important, determining whether the

relevant portion of the beam is successfully loaded into the moving magnetic trap - a rough estimate of this can be determined from the estimated beam velocity, but as shown in sections 4.1.1 and 4.1.3 detailed simulations are necessary in order to optimise this parameter. For argon the delay values are in the region of  $\approx 700$  to  $900 \mu\text{s}$ .

The wire quadrupole had been decommissioned when the data presented in this thesis were acquired; it is, however, included in some simulations of the experiment presented in chapter 4. When used, the wire quadrupole typically operated at  $\approx 40 \text{ V}$  and  $\approx 700 \text{ A}$ , see McArd [2].

Lastly, the detector used was an MCP, described in section 2.4.2. This has a maximum safe operating voltage of  $1.9 \text{ kV}$ , and we operated it at  $1.8 \text{ kV}$ .

## 2.6 Chapter Summary

In summary, this chapter has described the components of our Travelling-Wave Zeeman Decelerator experiment, including the source (sections 2.2.2 and 2.2.3), deceleration coils (section 2.3.4), power electronics (section 2.3.6), and detector (section 2.4.2). Key differences between our decelerator and similar experiments were discussed in sections 2.3.2 and 2.3.3. There were brief reviews of the operating principles of supersonic expansion sources (section 2.2.1), magnetic trapping and Zeeman deceleration (section 2.3.1), and time-of-flight signals (section 2.4.1).

Section 2.5 described the adjustable parameters of the experiment, and gave values for those used in the data presented in the following chapters. Some of these parameters, such as the DBD delay or valve pulse width, must be manually adjusted by trial and error for each run. Choosing an appropriate velocity for the travelling wave, however, relies on fitting a model to the detected signal, so as to infer the flow velocity and translational temperature of the molecular beam. The beam can then be simulated in order to assess how well the expected results match the data. The effective loading of the beam into a moving decelerator trap is largely determined by the power electronics delay, and simulations of the experiment were also necessary in order to optimise this parameter.

Chapter 3 describes the decelerator simulation codes, with examples of their application in reproducing results from the guiding and deceleration of metastable argon. Chapter 4 then describes the use of the simulation codes to determine optimised experimental parameters and make predictions about the results we would expect if the apparatus was modified, or different species were decelerated.

# Chapter 3

## Data Analysis and Simulation of the Zeeman Decelerator

Our travelling-wave Zeeman decelerator has been characterised using  $^3\text{P}_2$  metastable argon (see section 3.10). In order to interpret the resulting data a suite of computer models have been developed, which will be described in this chapter. Specific samples of data and sets of experimental parameters will be used as examples, with overall results given in a later section. When simulating a molecular beam we must begin with modelling the source, which would therefore be the logical place to begin this chapter; however, we rely on an analysis of the beam time-of-flight produced by the detector to provide the required parameters, so that will be the focus of the first section.

### 3.1 Analysing Time-of-Flight Data

#### 3.1.1 ToF at a Detector

Section 3.2 describes how we simulate the supersonic expansion source. The velocity profile of the initial state is estimated by fitting experimental data with a model of the expected signal (section 2.4.1), which will be covered here. The basic form of the signal is given by:

$$S(t) = B \frac{\Delta_{\text{SD}}^3}{t^4} \exp\left(-\left(\frac{\Delta_{\text{SD}}}{\sqrt{2}\sigma_{\text{T}}t_{\text{T}}}\right)^2 \left(\frac{t-t_{\text{T}}}{t}\right)^2\right), \quad (3.1.1)$$

which is equivalent to equation 2.4.9 in section 2.4.1. However, this expression does not include any broadening from the temporal and spatial widths of the beam source. In order to incorporate this we took an approach based on the model given in Tarbutt and Hinds [154]; if each particle is created at time  $t_0$  and position  $z_0$  then:

$$S(t, t_0, z_0) = B \frac{(\Delta_{\text{SD}} - z_0)^3}{(t - t_0)^4} \exp\left(-\left(\frac{\Delta_{\text{SD}} - z_0}{\sqrt{2}\sigma_{\text{T}}t_{\text{T}}}\right)^2 \left(\frac{(t - t_0) - (t_{\text{T}} - z_0/v_{\text{T}})}{t - t_0}\right)^2\right). \quad (3.1.2)$$

$v_{\text{T}}$  is the beam flow velocity. This is simplified by assuming that  $z_0 \ll \Delta_{\text{SD}}$ ,  $t_0 \ll t_{\text{T}}$ , setting  $t = t_{\text{T}}$  in the denominators of the two fractions and defining the width  $w = \sigma_{\text{T}}t_{\text{T}}^2/\Delta_{\text{SD}}$  and time  $t_{\rho} = -z_0/v_{\text{T}}$ , so that we can write:

$$S(t, t_0, t_{\rho}) = B \frac{\Delta_{\text{SD}}^3}{t_{\text{T}}^4} \exp\left(-\frac{(t - t_0 - t_{\text{T}} - t_{\rho})^2}{2w^2}\right). \quad (3.1.3)$$

If the terms  $t_0$  and  $t_\rho$  have the Gaussian distributions:

$$g(t_0) = \frac{1}{\sqrt{2\pi}\sigma_{t_0}} \exp\left(-\frac{(t_0 - \mu_{t_0})^2}{2\sigma_{t_0}^2}\right), \quad (3.1.4)$$

and:

$$h(t_\rho) = \frac{1}{\sqrt{2\pi}\sigma_{t_\rho}} \exp\left(-\frac{(t_\rho - \mu_{t_\rho})^2}{2\sigma_{t_\rho}^2}\right), \quad (3.1.5)$$

then the broadened signal at the detector is given by:

$$S'(t) = \int \int S(t, t_0, t_\rho) g(t_0) h(t_\rho) dt_0 dt_\rho, \quad (3.1.6)$$

which is given [154] as being of the form:

$$S'(t) = B \frac{\Delta_{\text{SD}}^3}{t_{\text{T}}^4} \frac{w}{w'} \exp\left(-\frac{(t - t_{\text{T}})^2}{2w'^2}\right), \quad (3.1.7)$$

where the new width  $w'^2 = w^2 + \sigma_{t_0}^2 + \sigma_{t_\rho}^2$ . (This assumes that any effects due to the response time of the detector electronics are negligible, but in principle extra terms could be added as necessary [175]). A problem with this model is that the source widths  $\sigma_{t_0}$  and  $\sigma_{t_\rho}$  are not known. We can, however, attempt to constrain them with the known properties of the DBD (section 2.2.3); if the DBD discharge time is  $t_{\text{D}}$  then  $\mu_{t_0} \approx t_{\text{D}}/2$  and  $\sigma_{t_0} \approx t_{\text{D}}/6$ . An average value for  $t_\rho$  is given by  $t_\rho \approx \Delta_{\text{D}}/2v_{\text{T}}$  where  $\Delta_{\text{D}}$  is the length of the excitation region, and similarly  $\sigma_{t_\rho} \approx \Delta_{\text{D}}/6v_{\text{T}}$  (for further consideration of these parameters see section 3.2).

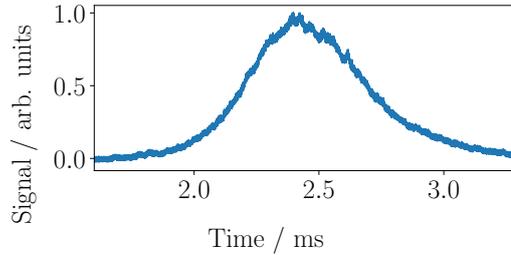


Figure 3.1: Example time-of-flight data for argon expansion; each data set is a time-average of 40 individual waveforms (see section 2.2.3 for discussion of the shot-to-shot stability of the source). By inspection approximate values of the peak time and overall width can be obtained; combining these with the known characteristics of the DBD excitation pulse we can then very roughly estimate the spatial, temporal and thermal width of the source.

A typical value for  $t_{\text{D}}$  is  $t_{\text{D}} \approx 30 \mu\text{s}$  giving  $\sigma_{t_0} \approx 5 \mu\text{s}$ . We know that  $\Delta_{\text{D}} = 12 \text{ mm}$ , but in order to estimate  $\sigma_{t_\rho}$  we also need a value of  $v_{\text{T}}$ , for which we need to extract  $t_{\text{T}}$  from a fit of equation 3.1.7. A very rough value can be found by inspecting a time-of-flight plot (figure 3.1) and noting that the central peak is at  $t \approx 2.4 \text{ ms}$ , so that  $v_{\text{T}} \approx \Delta_{\text{SD}}/2.4 \text{ ms} \approx 385 \text{ m s}^{-1}$ . This gives  $\sigma_{t_\rho} \approx 5 \mu\text{s}$ . Also by inspection, the bulk of the atoms appear to arrive within an interval of about 1.5 ms, so that  $w'$  is likely to be approximately  $250 \mu\text{s}$  (due to the  $3\sigma$  rule). These numbers suggest that the difference between  $w$  and  $w'$  should only be of the order of  $0.1 \mu\text{s}$ , and that satisfactory results may be achieved

by neglecting  $\sigma_{t_0}$  and  $\sigma_{t_p}$ , giving the simplified model:

$$S(t) = B \frac{\Delta_{SD}^3}{t_T^4} \exp\left(-\left(\frac{\Delta_{SD}}{\sqrt{2}\sigma_T t_T}\right)^2 \left(\frac{t-t_T}{t_T}\right)^2\right). \quad (3.1.8)$$

Figure 3.2 shows a fit of equation 3.1.8 to the data in figure 3.1. The fitting code initially trims the first 1 ms from the data to remove the ion peak (see section 2.4.3). It then determines the approximate time at which the signal has its maximum value, and that at which it falls to half of that value, giving an estimate of the half-width half-maximum (HWHM). The function  $S(t)$  is fitted to the signal data in the region extending  $\pm 5 \times \text{HWHM}$  from the centre (chosen so the fit covers the whole signal peak), for the parameters  $B$ ,  $t_T$  and  $\sigma_T$ . The fit is performed using a non-linear least-squares [178] function from the SciPy ‘optimize’ library [179]. The fitted values of  $t_T$  and  $\sigma_T$  give flow velocity  $v_T = 378 \text{ m s}^{-1}$  and translational temperature  $T_T = 6.4 \text{ K}$  (neglecting overall uncertainties in the quoted values, which will be addressed later). Repeating this process for equation 3.1.7, using the approximate values of  $\sigma_{t_0}$  and  $\sigma_{t_p}$  given above, produced estimates of  $T_T$  and  $v_T$  differing by less than  $0.1 \text{ K} / 0.1 \text{ m s}^{-1}$ . In addition to broadening the signal, the source widths should in theory shift the mean position. An alternative method to try and account for this is to adjust the arrival times in the data by subtracting the mean time of the DBD discharge  $t_D/2$ , and likewise reducing the source-detector distance  $\Delta_{SD}$  by  $\Delta_D/2$  in the fitting code. These effects, of shortening both the flight times and distances, partially cancel each other out, and introducing them was found to change the results negligibly.

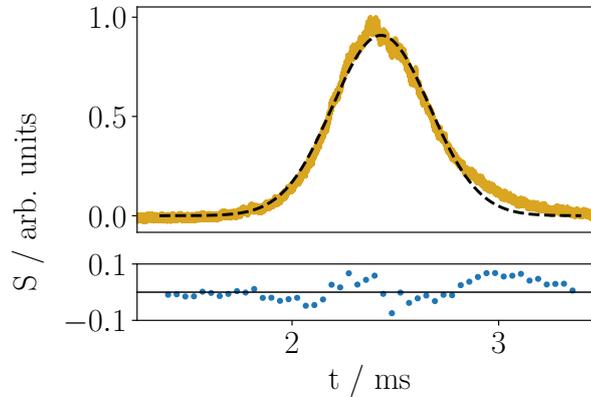


Figure 3.2: Least-squares fit of equation 3.1.8 (dashed line) to detector signal from argon supersonic expansion (orange), with stagnation pressure / temperature of 6 bar / 137 K. Estimated flow velocity of  $378 \text{ m s}^{-1}$  and translational temperature of 6.4 K were calculated from the fit (neglecting sources of uncertainty, which will be considered later). The deviation of the data from a purely Gaussian shape is obvious by inspection, and confirmed by the residual plot in the lower panel which shows significant structure.

Although they give consistent results with each other, both 3.1.7 and 3.1.8 are Gaussian functions due to the simplifications made in their derivation, and it can be clearly seen that the data depart from this. In fact, as will be described in section 3.2 we have no reason to expect that the data will assume such a shape; even an idealised velocity distribution based on a moving one-dimensional Maxwell-Boltzmann function will not generate a purely Gaussian time-of-flight. The apparent close

fit in the high-velocity / low-time tail of figure 3.2 is misleading, since if the expansion was occurring in accordance with the theory we should not see this; such an excessive ‘hot tail’ is seen in some other supersonic beam experiments [154]. Since we have argued that the effects of the width of the source are small for our experiment, we can plausibly attempt to apply equation 3.1.1. Figure 3.3 shows the latter fitted to the same data as the previous figures, and it can be seen that it appears to give a fit that is somewhat better in the low-velocity tail. This model gives  $v_T = 368 \text{ m s}^{-1}$  and  $T_T = 6.7 \text{ K}$ .

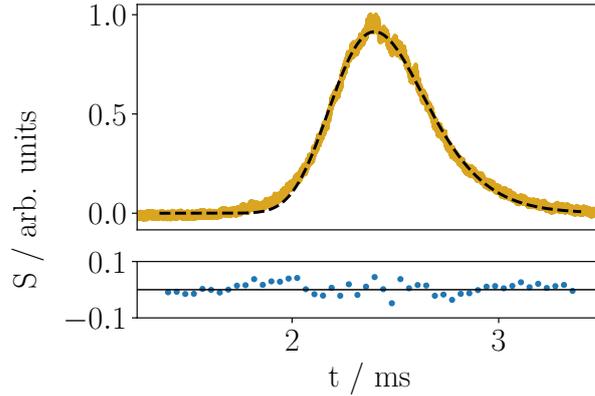


Figure 3.3: Least-squares fit of equation 3.1.1 (dashed line) to detector signal from argon supersonic expansion (orange), with stagnation pressure / temperature of 6 bar / 137 K. Estimated flow velocity of  $368 \text{ m s}^{-1}$  and translational temperature of 6.7 K were calculated from the fit. Residuals plotted in the lower panel show a closer fit than equation 3.1.8, although there is still structure.

Further insight into data fitting can be obtained by subtracting the model from the data to give residuals, as shown in the lower panels of figures 3.2 and 3.3; significant structure can be seen in both. We attempted to smooth the data to reduce the spread of values about the mean along the length of the trace, in order to determine whether this would improve the nature of the fit. As described in section 2.2.3, a time-of-flight trace as shown in figure 3.1 is actually produced by taking an average of 40 individual traces from the experiment, in order to counter shot-to-shot instability in the source and detector. The effects of outliers on the standard error of a fit can be further reduced by smoothing the data.

Two techniques were used; firstly the data were re-binned. For  $N$  data points of the function  $f(x)$ , re-binning with parameter  $M$  splits the data into sets of  $M$  points and takes the average of  $f(x)$  in each interval, shrinking the size of the data set to  $N/M$ . Secondly, a moving-average was used, in which the value of  $f(x)$  at each point  $x$  is replaced with the mean of the values  $f$  in the interval  $x \pm M/2$  (this results in shrinking the number of data points by  $M$ ). Figure 3.4 shows the example data used previously in this chapter adjusted with a moving average, parameter  $M = 2 \times 10^4$ , and then fitted with equation 3.1.1. The trace is visibly smoothed, with the estimated values of  $v_T / T_T$  altered by  $\approx 1\% / 10\%$ , but although the spread in the residuals is slightly reduced the structure remains. Similar results were found for simple re-binning, and in both cases  $M$  had to be at least of order  $10^4$  to make a detectable difference to the fit. Setting such large values of  $M$  risks distorting the shape of the signal for no real benefit, and so neither approach was pursued.

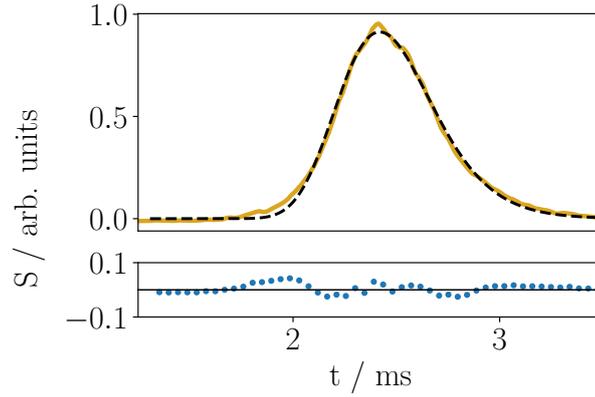


Figure 3.4: Fit of equation 3.1.1 (dashed line) to data as per figure 3.3, but data smoothed with a moving-average,  $M$  parameter of  $2 \times 10^4$ . Estimated  $v_T / T_T$  are altered by 1% / 10% but the residuals are not significantly improved.

To summarise; we have two models for our data (or, two variants of one underlying model) both of which give a reasonable match to real data, from which we can infer plausible values for the beam velocity and temperature. There is significant structure to the residuals of both models, with the same overall shape in either case, indicating that they are only approximate representations of the actual distribution of energies in the beam. This is perhaps not surprising; we should not expect too close a match between theory and data, given the number of simplifying assumptions built into all of the models of the expansion that have been discussed. The two sets of values for  $v_T$  and  $T_T$  also differ significantly from each other, by roughly 5 % and 6 %, which leaves us two issues to contend with; which model gives the best representation of the experiment, and how do we quantify the errors in the values that it gives us?

Regarding the latter, further analysing the model fit by considering the covariance and correlation values of the fitted parameters is not meaningful, since the errors on the fit cannot tell us anything about the uncertainties in the beam characteristics. Regardless of the model, these uncertainties are defined by the spatial and temporal characteristics of the experiment. The MCP has a time response on the order of 1.5 ns (section 2.4.2), and the Even-Lavie valve opens with a delay of a few  $\mu\text{s}$  (section 2.2.2). (These are hard to quantify, including their linearity, but they should be small enough to neglect). Additional delays introduced by the electronics should also be small in comparison to the beam time-of-flight profile, which is of the order of several ms. The error is likely to be dominated by uncertainties in our knowledge of the distance from the source to the detector,  $\Delta_{SD}$ . The width of the signal peak determines the beam temperature (or, equivalently, its spread of velocities) and this is dependent on both time and distance factors, namely the parameters  $\sigma_T$ ,  $t_T$  and  $\Delta_{SD}$ . Again, errors in  $\Delta_{SD}$  will be dominant. Choosing a value for the uncertainty in  $\Delta_{SD}$  of 5 mm (and it is likely to be smaller than this) leads to an error in  $v_T$  of  $\approx \pm 3 \text{ m s}^{-1}$ , and in  $T_T$  of  $\approx \pm 0.1 \text{ K}$ . Although this analysis is not particularly rigorous it is probably appropriate, given the apparent mismatch between theory and experiment shown in the structured residuals of the fits. Bearing in mind the adage that ‘all models are wrong, but some of them are useful’ [180], we return to the question of which approach gives the best match to our data. To decide this, Monte Carlo simulations of the generation, propagation and detection of the beam were developed, based on the parameters extracted from equations 3.1.8 and 3.1.1. These simulations are discussed in section 3.2.

### 3.1.2 ToF at Dual Detectors

Section 3.1.1 describes fitting a model of the expected time-of-flight signal arising from a supersonic expansion to experimental data, with the goal of extracting an estimate of the beam velocity distribution. Equation 3.1.7 gave a form for the signal:

$$S'(t) = B \frac{\Delta_{SD}^3}{t_T^4} \frac{w}{w'} \exp\left(-\frac{(t-t_T)^2}{2w'^2}\right), \quad (3.1.7)$$

with the width  $w'^2 = w^2 + \sigma_{t_0}^2 + \sigma_{t_\rho}^2$  being composed of thermal, temporal and spatial terms. We previously argued that the latter two are negligible for our experiment but a method of eliminating them is described in Tarbutt and Hinds [154]. This requires sampling the beam with two spatially-separated detectors, giving  $w_2'^2 - w_1'^2 = w_2^2 - w_1^2$ . The thermal broadening is given by  $w = \sigma_T t_T / v_T$ , so that the translational temperature  $T_T$  can be determined:

$$T_T = \frac{w_2'^2 - w_1'^2}{t_{T2}^2 - t_{T1}^2} \frac{m v_T^2}{k_B}. \quad (3.1.9)$$

In previous work on the decelerator project a bespoke fast-ionisation gauge (FIG) was used [1]. This uses electrons produced by thermal emission from a hot filament to ionise some of the atoms in the beam, which are accelerated to a collecting electrode by an electric field, producing a measurable signal. The FIG was reinstalled between the Even-Lavie valve and skimmer (see figure 2.2). A supersonic expansion of argon was then produced and the signals at both the FIG and MCP captured. Both sets of data were fitted with Gaussian functions to determine estimates of the standard deviation  $w'$  and central time  $t_T$ . The beam flow velocity  $v_T$  was estimated from the known distance of the MCP from the source  $\Delta_{SD}$  and the arrival time  $t_{T2}$ . An example is shown in figure 3.5.

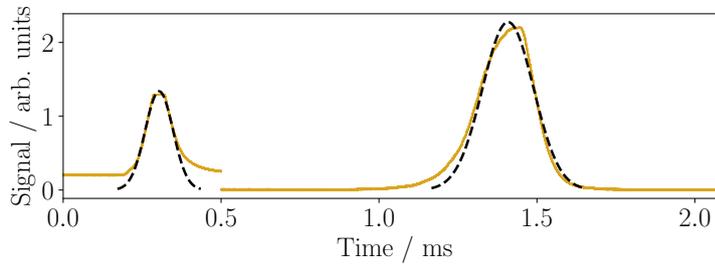


Figure 3.5: Time-of-flight data for supersonic argon expansion, captured on fast-ionisation gauge (left) and microchannel-plate detector (right). In theory these can be fitted with Gaussian functions and combined to eliminate unknown spatial and temporal widths of the source to extract an estimate of the translational temperature of the beam. In practice the results were inconclusive, due to the different operating principles of the devices.

Data were taken with the valve at room temperature for various sets of parameters for the valve, DBD and FIG. The calculated temperatures were compared with those given by fits of equation 3.1.8, and it was found that the two methods give estimates of the temperature that differ significantly. However; the two detectors operate on different principles, so their outputs do not necessarily represent equivalent parts of the actual particle distribution. The FIG also causes heating of the sample of atoms. It was concluded that this method is not applicable to our experimental layout. As described in section 3.2, the methods we have used to estimate the beam parameters from the MCP detector alone appear to give results that are accurate enough for our simulation requirements. Alternatively, we could utilise this method by remounting the MCP on a movable stage.

## 3.2 Supersonic Expansion Source

This section describes the generation of the initial conditions of the atoms in our simulations. Before the interactions with the decelerator fields are included it is important to establish that we can reliably reproduce our zero-field ‘transmission’ data, such as the time-of-flight profile shown in figure 3.1. There are several difficulties involved in modelling our beam source. Simulating the actual redistribution of particle energies that produces the beam from the population in the stagnation chamber would be a significant task, involving a consideration of collisional cross-sections, and many-body collisions are very computationally expensive to calculate. An option is to produce our initial ensemble of atoms a short distance from the valve nozzle, assuming that the expansion is complete; this was initially done as a one-dimensional case.

The simplest model would treat the atoms as a point source emerging in negligible time, but we do have an idea of the spatial and temporal widths of the valve, and we can in confidence at least define their probable mean values. We do not know the actual distribution functions for these parameters but we can consider them as being Gaussian [154] (see section 3.1.1). Defining the front of the valve housing as being at  $z = 0$ , the nozzle is at  $z = -6$  mm, and the initial positions  $z_0$  are given by:

$$h(z_0) = \frac{1}{\sqrt{2\pi}\sigma_{z_0}} \exp\left(-\frac{(z_0 - \mu_{z_0})^2}{2\sigma_{z_0}^2}\right), \quad (3.2.1)$$

with mean  $\mu_{z_0} \approx -5$  mm and standard deviation  $\sigma_{z_0} \approx 1$  mm (this is equivalent to equation 3.1.3). We have mentioned that we are assuming that the isentropic expansion has effectively ceased and the beam temperature profile has stabilised within a short distance from the nozzle, however the collision rate of a supersonic expansion typically only becomes negligible after several tens of nozzle diameters  $D$  from the source [49]. The Even-Lavie valve has a 0.2 mm aperture, so that the estimated parameters above are not ideal, although the expansion is very likely to be complete by the time the beam pulse passes the front face of the valve. Similarly, we assume that the creation times  $t_0$  of the atoms in the beam are:

$$g(t_0) = \frac{1}{\sqrt{2\pi}\sigma_{t_0}} \exp\left(-\frac{(t_0 - \mu_{t_0})^2}{2\sigma_{t_0}^2}\right), \quad (3.1.4)$$

with  $\mu_{t_0}$  and  $\sigma_{t_0}$  defined by the overall opening time of the valve  $t_v$ , typically between 20 and 30  $\mu\text{s}$ , and it takes  $\approx 10$   $\mu\text{s}$  for the plunger to open after activation. The mean creation time is therefore  $\mu_{t_0} = (10 + t_v/2)$   $\mu\text{s}$ , with an estimated  $\sigma_{t_0} \approx t_v/6$ . Supersonic expansions are often specified as a velocity distribution with different components in the directions parallel to and orthogonal to the beam flow [49]. This is shown by equation 2.2.5 in section 2.2.1, the longitudinal component of which is:

$$f_{\text{T}}(v_z) = \left(\frac{1}{2\pi\sigma_{\text{T},z}^2}\right)^{1/2} \exp\left(-\frac{(v_z - v_{\text{T}})^2}{2\sigma_{\text{T},z}^2}\right). \quad (3.2.2)$$

A range of initial particle velocities can be chosen by sampling from equation 3.2.2; the required parameters are found by analysis of experimental data, as described in section 3.1.1. Putting all of this together, we can create our sample of atoms and evolve their motion through space. Since the three initial conditions are all Gaussian functions they can be efficiently sampled via the NumPy ‘normal’ library [181]. We then run a stepwise simulation, with each particle only beginning to move once the

time counter has passed its creation time. For the transmission model there are no forces involved, so the displacement of each particle from its position  $z_n$  at time  $t_n$  is given by  $z_{n+1} = z_n + v_z \delta_t$ , where  $v_z$  is its velocity and the timestep  $\delta_t = t_{n+1} - t_n$ .

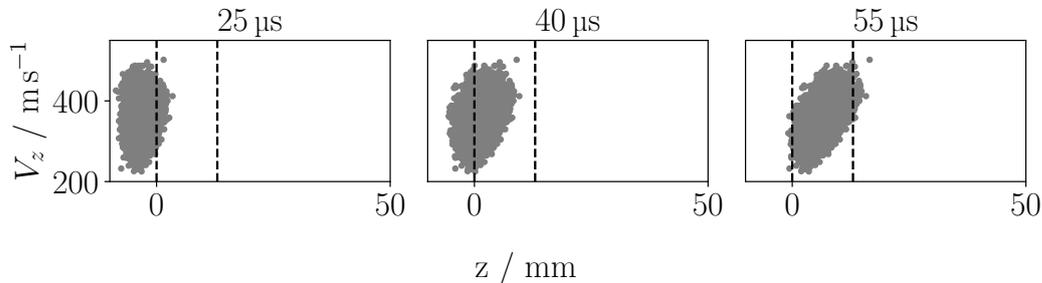


Figure 3.6: Simple one-dimensional model for expansion of argon atoms, the valve opens at time  $t = 0$  for  $25 \mu\text{s}$ ,  $v_T = 361 \text{ m s}^{-1}$ , temperature of  $6.9 \text{ K}$  giving  $\sigma_{T,z} = 37.8 \text{ m s}^{-1}$ . Dashed lines indicate the dimensions of the DBD which runs to  $z = 12 \text{ mm}$ . The difficulty of evenly exposing the entire expansion to an excitation pulse lasting  $30 \mu\text{s}$  can be seen.

Figure 3.6 shows an example of this simulation applied to argon. In the case of argon we are using a dielectric barrier discharge pulse (DBD, section 2.2.3) to produce metastable atoms, so that they are both paramagnetic and therefore manipulable by the decelerator field, and also possess sufficient internal energy to trigger the microchannel-plate detector (MCP, section 2.4.2). The use of metastable argon (or some other species that must be energised to an excited state), has implications for our implicit assumption that the parameters estimated from the signal at the detector match the underlying distribution of the beam as it emerges from the valve. We are only able to detect excited atoms, and it is difficult to quantify the differences between the velocity distributions of these two populations. There are two primary reasons that these may diverge; firstly, as well as producing the desired excited atoms, the DBD pulse has the side-effect of causing heating due to ionisation of trace impurities in the beam. These effects are not as serious for a DBD as compared to traditional electron-impact excitation schemes [158], but they are present.

Secondly, we generally wish to optimise the duration of the DBD discharge, given by the number of pulses  $N$  divided by the RF frequency,  $\nu$ , and the time  $t_D$  at which it is triggered, in order to select a slower fraction of the beam and increase the chances of loading it into the decelerator. This means that we are certain to distort the distribution of particle velocities we infer from the detected time-of-flight signal. In fact, even if we wish to uniformly excite the entire beam this is difficult to achieve, as illustrated in figure 3.6. The packet of atoms expands along the  $z$ -axis as it moves, due to thermal spread of velocities, and rapidly becomes larger than the  $12 \text{ mm}$  DBD. A typical value for  $N/\nu$  would be  $18 / 600 \text{ kHz} = 30 \mu\text{s}$ , and even if the majority of atoms are inside the region of the DBD within this time interval they will not be uniformly exposed. The fact that the discharge may not be uniform along the axis of the beam is also hard to quantify. Simply extending the duration of the discharge is not always feasible, as generally the values of  $N$  and  $\nu$  must be empirically set in order to maximise beam signal and stability. This argument is of course specifically for argon and may not be as applicable to heavier species with lower supersonic expansion flow velocities, but generally the point is that when using excitation on a beam we must expect the detected signal to differ from the one we would theoretically detect, in ways that may be hard to quantify. This

is likely to contribute to the differences we saw when fitting models to data in section 3.1.1. With the above caveats in mind, we proceed to simulate the source population as it is created during the excitation pulse. The DBD is 12 mm long so we can use equation 3.2.1 for the positions  $z_0$ , with  $\mu_{z_0} = 6$  mm and  $\sigma_{z_0} = 2$  mm. Choosing time  $t = 0$  at the DBD trigger, the creation times  $t_0$  are bounded  $0 \leq t_0 \leq N/\nu$ , giving  $\mu_{t_0} = N/2\nu$  and  $\sigma_{t_0} = N/6\nu$ . We obtain estimates of  $t_T$  and  $\sigma_T$  by fitting equation 3.1.8:

$$S(t) = B \frac{\Delta_{SD}^3}{t_T^4} \exp\left(-\left(\frac{\Delta_{SD}}{\sqrt{2}\sigma_T t_T}\right)^2 \left(\frac{t-t_T}{t_T}\right)^2\right), \quad (3.1.8)$$

these are applied to equation 3.2.2 (as  $v_T = \Delta_{SD}/t_T$ ) to create  $v_z$ . The motion of the atoms through space is then solved until they reach a distance corresponding to the detector, at which point their arrival times are recorded. The loop terminates when all are accounted for, and the array of arrival times can be plotted as a histogram and directly compared to the data from which their initial conditions were estimated.

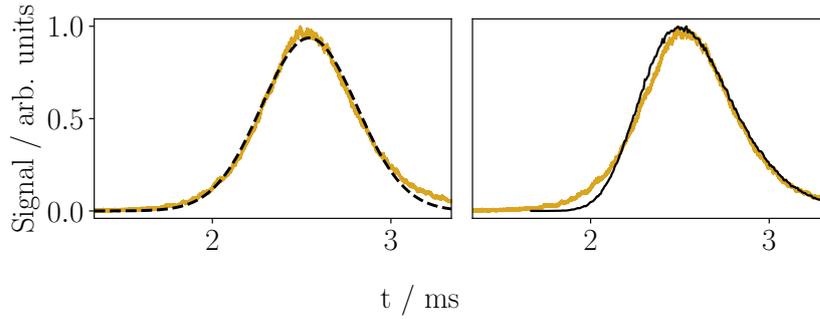


Figure 3.7: Argon supersonic expansion (orange), stagnation pressure / temperature of 6 bar / 137 K. Left, dashed line; fit of equation 3.1.8, giving estimated flow velocity  $v_T = 361$  m s<sup>-1</sup> and temperature  $T_T = 6.8$  K. Right, black line; simulated time-of-flight from a Gaussian velocity source using these parameters.

As stated in section 3.1.1, equation 3.1.8 is problematic; a Gaussian velocity distribution does not give rise to a Gaussian time-of-flight profile, as confirmed by figure 3.7. The right plot shows that the signal should diverge from a pure Gaussian in both the low- and high-velocity tails. The apparent good fit of the data in the high-velocity tail as shown in the left plot indicates a divergence of the data from theory. However, the match at low velocities is reasonable, and this is the regime we are most interested in, since we have more chance of manipulating slower particles with the decelerator. Tarbutt *et al.* [154] also reported an additional population of fast atoms when fitting their model of a supersonic expansion for various monatomic carrier gases; they state that optimising their source for low temperature rather than maximised signal and stability improved the fit in this region. As described in section 2.2.3 we experienced low shot-to-shot stability and low signal when using argon in our Even-Lavie valve, which was improved with the installation of a thermal electron filament, which may be relevant. The degree to which our data differ from predictions is also variable; see for example the fits in section 3.1.1, which are for different data than the ones in this section. The data show a noticeably different population in the high-velocity tail, and during an experimental run the characteristics of the beam will drift due to uncontrollable ambient conditions.

Lastly, we investigated the effects of the assumed distributions given in equations 3.2.1 and 3.1.4. Versions of this simulation were produced in which the initial conditions  $t_0$  and  $z_0$  were either all set to their mean values, or uniformly (as opposed to normally) distributed within their bounds. The output of these tests could not be distinguished from simulations in which the parameters were sampled from equations 3.2.1 and 3.1.4, which supports our claim in section 3.1.1 that the temporal and spatial widths of the source are negligible in our experiment.

The simulations described up to this point are one-dimensional, and were later extended to include the effects of the decelerator field, as discussed in a later section. For that case it was more important to include transverse motion in the beam, given the strongly directional profile of the field. Three-dimensional extensions to the model were first developed and tested for free flight, which will now be discussed. The radial component  $v_r$  of the source velocity in equation 2.2.5 is:

$$f_{\Gamma}(v_r) = \frac{1}{2\pi\sigma_{\Gamma,r}^2} \exp\left(-\frac{v_r^2}{2\sigma_{\Gamma,r}^2}\right), \quad (3.2.3)$$

however, we do not know what the values of the parameters  $v_r$  and  $\sigma_{\Gamma,r}$  are. When dealing with the unknown longitudinal velocity components we are able to extract estimates for the relevant parameters from the time-of-flight data, but our detector does not allow us to determine anything about the radial velocities. We can, however, use the fact that the beam passes through a skimmer to determine plausible initial conditions for each atom, since this constrains the possible values of the coordinates  $x$ ,  $y$  and velocity components  $v_x$ ,  $v_y$ . For each particle a trajectory is defined passing from a point at the valve to a point at the skimmer, with radial and angular coordinates  $d_V, \theta_V$  and  $d_S, \theta_S$ . These are chosen from a uniform random distribution, with the radial values bounded by the radii of the valve / skimmer, which are 0.1 mm and 2 mm respectively. These coordinates are easily put into Cartesian form:

$$\begin{pmatrix} x_V \\ y_V \\ x_S \\ y_S \end{pmatrix} = \begin{pmatrix} d_V \cos(\theta_V) \\ d_V \sin(\theta_V) \\ d_S \cos(\theta_S) \\ d_S \sin(\theta_S) \end{pmatrix}. \quad (3.2.4)$$

Although the trajectories start at the valve, we create the particles at  $z = z_0$ , so the initial lateral coordinates  $x_0, y_0$  are determined as being on the line joining the initial / final points  $x_V, y_V / x_S, y_S$ :

$$\begin{pmatrix} x_0 \\ y_0 \end{pmatrix} = \begin{pmatrix} \frac{(x_S - x_V)}{\Delta_{VS}} z_0 + x_V \\ \frac{(y_S - y_V)}{\Delta_{VS}} z_0 + y_V \end{pmatrix}, \quad (3.2.5)$$

where  $\Delta_{VS} = 162$  mm is the distance from the valve to the skimmer. Each particle has to travel a distance  $\Delta_z = \Delta_{VS} - z_0$ , which takes time  $t_S = \Delta_z / v_z$ . In this time each particle moves  $\Delta_x = x_S - x_0$  and  $\Delta_y = y_S - y_0$  so the initial transverse velocity components are given by:

$$\begin{pmatrix} v_x \\ v_y \end{pmatrix} = \begin{pmatrix} \frac{\Delta_x v_z}{\Delta_z} \\ \frac{\Delta_y v_z}{\Delta_z} \end{pmatrix}. \quad (3.2.6)$$

To create the initial distribution in a three-dimensional source model, the polar coordinates are selected from the uniform random ranges  $0 \leq d_V \leq 0.1$  mm,  $0 \leq d_S \leq 2$  mm,  $0 \leq \theta_V \leq 2\pi$ ,  $0 \leq \theta_S \leq 2\pi$ ; these are used to calculate the components  $x_0, y_0, v_x, v_y$ . Values for  $z_0, v_z$  and  $t_0$  are chosen as described earlier.

Assuming, as we have done, that the radial velocities are independent of the longitudinal ones, their presence would make no difference to the expected time-of-flight signal except for the complicating factor that in the physical experiment the beam travels from the skimmer to the detector through a tube with an inner radius of 2.62 mm. To account for this the three-dimensional simulations count particles that reach this radial distance during flight as being lost due to collision with the tube wall. This means that on average particles with higher longitudinal velocities can have a slightly larger range of possible radial velocities, since they have less time to drift to the wall during flight. For a long enough decelerator this would shift the signal at the detector, with the slower tail of the distribution being winnowed out, but in practice this effect turns out to be negligible for the typical beam speeds we achieve with metastable argon. The distribution of longitudinal velocity components of atoms ‘surviving’ the trip to the detector in the three-dimensional simulation differ from those of the one-dimensional model in no statistically measurable way. The effect may, however, be more significant if slower beams of different species were used.

Returning to the velocities of the particles; so far we have used equation 3.2.2 with parameters from 3.1.8, both of which are Gaussian, to generate a longitudinal velocity distribution. In section 3.1.1 we also demonstrated a fit of equation 3.1.1, which gives an arguably closer match to the data, particularly in the low-velocity wing in which we are most interested:

$$S(t) = B \frac{\Delta_{SD}^3}{t^4} \exp\left(-\left(\frac{\Delta_{SD}}{\sqrt{2}\sigma_T t_T}\right)^2 \left(\frac{t-t_T}{t}\right)^2\right). \quad (3.1.1)$$

The deviation of this function from a ‘pure’ Gaussian shape means that we cannot use it to generate inputs for the typical velocity distribution given in equation 3.2.2, but we can use it as the basis for an alternative way to model the source. Returning to section 2.4.1; equality 2.4.8 related the time-dependent signal  $S(t)$  to  $\Phi_T(v)$ , the flux of particles as a function of speed:

$$S(t) \propto \Phi_T(v)/t. \quad (2.4.8)$$

Equality 2.4.8 shows that the beam speed distribution can be retrieved directly from the time-of-flight data, since a suitably-normalised plot of  $S(t)t$  against  $1/t$  is equivalent to one of  $\Phi_T(v)$  against  $v$  [173] (assuming, as we have done, that broadening of the signal due to the width of the source is small). We can define the expected signal as a function of velocity  $S(v) \equiv \Phi_T \equiv S(t)t$ , giving:

$$S(v) \equiv S(t)t = Bv^3 \exp\left(-\frac{(v-v_T)^2}{2\sigma_T^2}\right). \quad (3.2.7)$$

Figure 3.8 shows the example data used in this section converted to a speed distribution by plotting the product of the signal and time against  $v = \Delta_{SD}/t$ , and fitted with equation 3.2.7 for the parameters  $B, v_T$  and  $\sigma_T$ . This gives the same results as equation 3.1.1, but this form is convenient for simulation purposes. We create our initial values  $x_0, y_0, z_0$  and  $t_0$  as described earlier in this section, but now instead of velocity components set by the one-dimensional Maxwell-Boltzmann distribution

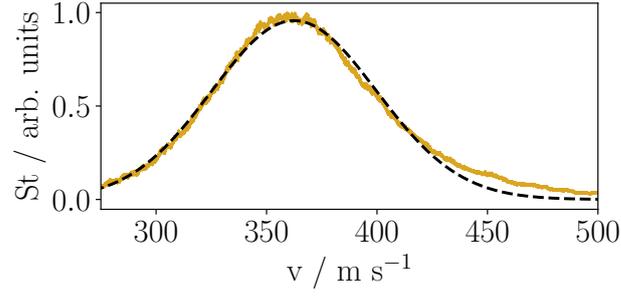


Figure 3.8: Orange, speed distribution derived from experimental time-of-flight data, since  $S(t)t$  vs.  $1/t \equiv \Phi_T(v)$  vs.  $v$ . Dashed line, fit of equation 3.2.7. The results of this are equivalent to fits of equation 3.1.1 as shown in e.g. figure 3.9.

and the geometry of the valve and skimmer, we choose initial speeds  $v_0$  from the distribution:

$$f(v_0) = Bv_0^3 \exp\left(-\frac{(v_0 - v_T)^2}{2\sigma_T^2}\right). \quad (3.2.8)$$

Unlike the Gaussian case we cannot directly sample this function using a NumPy library but instead use Monte Carlo rejection sampling to create a population of the desired size (see section 3.3). This is slower than sampling from equation 3.2.2 but it is still a small part of the overall runtime of the simulation. Each atom has already been assigned a random trajectory from the valve to the skimmer which can be written as a displacement vector  $\mathbf{r}$ , with magnitude  $r = |\mathbf{r}|$  and components:

$$\begin{pmatrix} r_x \\ r_y \\ r_z \end{pmatrix} = \begin{pmatrix} x_S - x_V \\ y_S - y_V \\ \Delta_{VS} \end{pmatrix}, \quad (3.2.9)$$

giving angles:

$$\begin{pmatrix} \theta_X \\ \theta_Y \\ \theta_Z \end{pmatrix} = \begin{pmatrix} \arccos(r_x/r) \\ \arccos(r_y/r) \\ \arccos(r_z/r) \end{pmatrix}, \quad (3.2.10)$$

from which we can resolve the speed  $v_0$  into vector components  $v_x$ ,  $v_y$  and  $v_z$ :

$$\begin{pmatrix} v_x \\ v_y \\ v_z \end{pmatrix} = v_0 \begin{pmatrix} \cos(\theta_X) \\ \cos(\theta_Y) \\ \cos(\theta_Z) \end{pmatrix}. \quad (3.2.11)$$

Once the initial coordinates are established the simulation can proceed as previously described. This is shown in figure 3.9, the right-hand pane of which has the simulated time-of-flight versus data. On the left is a fit of equation 3.1.1, which as stated is equivalent to equation 3.2.7.

Compared to figure 3.7 it can be seen that both methods give a very similar output for the simulated signal, but only for the speed model do we see an agreement between the model fitted to the data and the output of the simulation. Supersonic expansions involve complex physical interactions, and any model is bound to include many approximations. We have developed plausible models for the

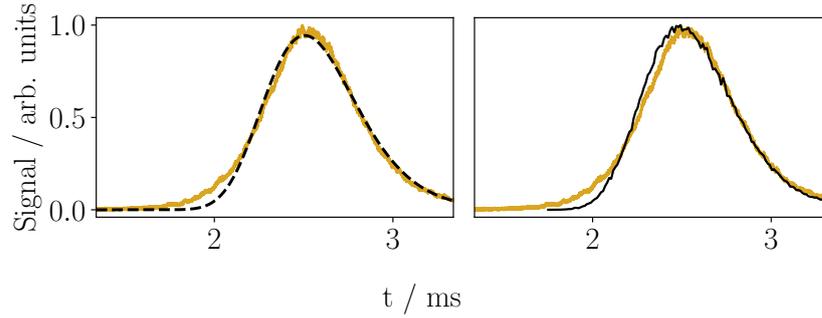


Figure 3.9: Argon supersonic expansion (orange), stagnation pressure / temperature of 6 bar / 137 K. Left, dashed line; fit of equation 3.1.1, giving estimated flow velocity  $v_T = 350 \text{ m s}^{-1}$  and temperature  $T_T = 6.9 \text{ K}$ . Right, black line; simulated time-of-flight with initial speeds selected from equation 3.2.8 and trajectories devised on the geometry of the valve / skimmer.

distribution of velocities in the molecular beam, with varying results, but the technique of sampling from an arbitrary function offers us another option. Rather than fitting some theory to data, we can just look at real time-of-flight profiles and use them as a template for the initial conditions of the beam in our simulations. We have fitted equations 3.1.8 and 3.2.7 to data using least-squares methods but instead we can apply a one-dimensional cubic spline (section 3.7.1); being a piecewise model this can reproduce the data closely along its entire range, without any propagation of local errors. This spline interpolator can be used with Monte Carlo rejection sampling to produce a range of initial speeds that very closely correspond to the data, without making any assumptions about the theoretical basis for their distribution; otherwise the initial conditions are determined as previously described. The results of such a simulation are shown in figure 3.10, and it is clear that this represents the best possible case for a simulation to match the source data.

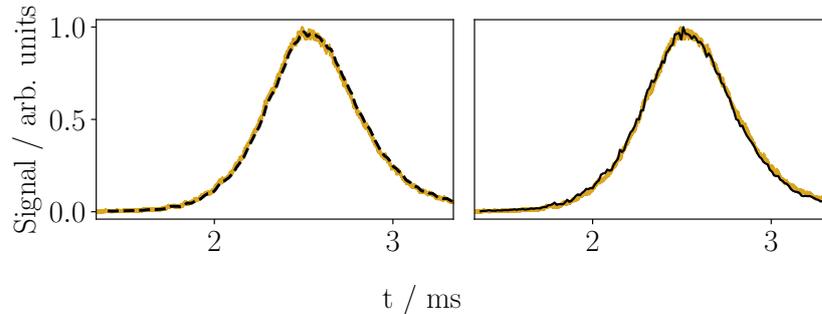


Figure 3.10: Argon supersonic expansion (orange), stagnation pressure / temperature of 6 bar / 137 K. Left, dashed line; fit of a one-dimensional cubic spline to experimental data. Right, black line; simulated time-of-flight with initial speeds selected by Monte Carlo rejection from the spline interpolator.

In conclusion; in section 3.1.1 we discussed the derivation of a commonly-used Gaussian model for the expected time-of-flight signal arising from a supersonic expansion, and compared it to a parent model with fewer simplifying assumptions. When used as inputs for a Monte Carlo simulation of the beam both methods plausibly reproduce the gross structure of the actual time-of-flight data, but we have achieved ideal results by simply fitting an ad-hoc model to zero-field transmission data. This serves as the basis for a simulation of the full decelerator, including magnetic fields.

### 3.3 Arbitrary Function Sampling

In Monte Carlo simulation it is often necessary to create a random sample of a variable  $x$  which is distributed according to a function  $f(x)$ . When  $f$  is a common function such as a Gaussian then this is straightforward, since most programming languages have standard libraries for creating such populations; for example, the NumPy ‘random’ library for Python [181]. These are typically easy to use and optimised for efficient and fast operation, although care needs to be taken that the behaviour of the random number generator (RNG) is as expected. Historically there were problems in the context of codes utilising multiprocessing, involving parallel processes all having the same value of some random parameter rather than unique ones. These issues have been largely resolved in modern libraries but care should be taken with legacy codesets as unintended behaviour may distort the results of simulations in ways that are difficult to detect. Even when operating as intended it must be borne in mind that all code libraries of this type are actually pseudo rather than true random number generators, and for applications with very strict requirements for ‘true’ randomness specialist hardware equipment is available. When using software generation it should be noted that not all algorithms are equal, and many commonly used codes actually fail statistical testing. For very sensitive applications it may be advisable to perform analysis to test the RNG to determine that the output of the code is not compromised. Recent versions of NumPy have moved away from the traditional Mersenne Twister 19937 method [182] and default to the statistically better Permuted Congruential Generator PGC64 [183]. This is a deep topic, but for the purposes of the work described in this thesis it is sufficient to note that all results given have utilised at least PGC64 where a pseudo-RNG was required, with legacy code being updated or rewritten completely.

If  $f$  is some arbitrary function then other methods must be used to randomly sample it. Depending on its form a method such as inverse transform sampling [184] may be a good solution, but finding the inverse of the cumulative distribution function is not always simple. There are many other options but one that is easy to apply and will always work is Monte Carlo Rejection sampling [185], although it is not necessarily very computationally efficient. The operation of this method is straightforward, requiring as input the maximum value of the function  $f_{\max}$  and the minimum and maximum values of the interval in which we wish to sample,  $x_{\min}$  and  $x_{\max}$ . To begin, two random numbers  $n_1$  and  $n_2$  in the range  $[0,1]$  are generated and used to give:

$$x = x_{\min} + n_1(x_{\max} - x_{\min}), \quad (3.3.1)$$

and:

$$y = n_2 f_{\max}. \quad (3.3.2)$$

If  $y \leq f(x)$  then we accept  $x$  as a random value sampled from the function  $f$ , otherwise we reject it. This process continues until the desired number of values  $x$  has been generated. The interval parameters  $x_{\min}$  and  $x_{\max}$  must be chosen appropriately such that they bound the function  $f$  across the entire region in which it has a substantial non-zero value, but this is a trade-off for longer computation time since expanding the interval of  $x$  will greatly increase the number of null results of the rejection tester. Also, the value of  $f_{\max}$  may or may not be simple to determine from the form of  $f$ ; for example, if  $f$  is a Gaussian it is clear that  $f_{\max} = \mu$ , the mean value of  $x$ . For an easily differen-

table function we can find the maximum value by solving  $\frac{d}{dx}f(x) = 0$ . An alternative method is to generate a range of values of  $x$  that bracket the peak of the function, calculate the corresponding values  $f(x)$  and then inspect the output array for the maximum value. The spacing step between neighbouring values of  $x$  in the query can be made small to allow an arbitrarily close value of  $f_{\max}$  to be found; again, depending on the form of  $f$  this may not be computationally efficient, but it may be simpler than finding the root of the derivative and will always work for a continuous function  $f(x)$ .

To give an example, in section 3.2 we discuss the generation of a range of particle speeds from the distribution shown in equation 3.2.8:

$$f(v_0) = Bv_0^3 \exp\left(-\frac{(v_0 - v_T)^2}{2\sigma_T^2}\right). \quad (3.2.8)$$

This function is a skewed Gaussian, so we can choose  $x_{\min} / x_{\max}$  to be  $v_T \pm 3\sigma_T$ , meaning that our sampled distribution should be drawn from better than 99 % of the population, thus being representative of the overall distribution. The maximum value of this function is easily found to be one of the roots:

$$f_{\max} = \frac{1}{2} (v_T \pm (v_T^2 + 12\sigma_T^2)). \quad (3.3.3)$$

As stated, this method is relatively slow compared to generation of purely Gaussian samples using NumPy; for example, generating a population of  $10^6$  values of  $v_0$  on a recent desktop computer took over 10 minutes to calculate compared to less than one minute for the Gaussian case. Since our overall Monte Carlo simulation of the experiment may take several hours to run, however, this remains a small part of the overall code runtime, and it only occurs once at the beginning of the code. The slowdown can be mitigated; because the workload of the algorithm can be broken into independent chunks it is easily parallelisable. If frequent calls to this method were needed during runtime then it would certainly be a priority to determine a more efficient method.

### 3.4 Magnetic Trapping

Section 2.3.1 gave a brief review of the Zeeman effect, and the associated Stern-Gerlach mechanism by which particles with a magnetic dipole moment may be manipulated using an inhomogeneous magnetic field; this section will describe how these forces may be calculated for various species. The travelling-wave Zeeman decelerator has been characterised with argon in the metastable  $^3P_2$  state so that will be considered first. This term has  $L = 1$ ,  $S = 1$ ,  $J = 2$  and  $I = 0$ , giving a simple structure with five equally-spaced sub-levels. Previously it was stated that for moderately strong fields the Zeeman splitting in ground-state silver took the form:

$$\Delta E \approx g_J \mu_B m_J B_z, \quad (2.3.8)$$

because  $I = 0$  this can be suitably applied to argon across the range of field strengths of interest in our experiment, as shown in figure 3.11.

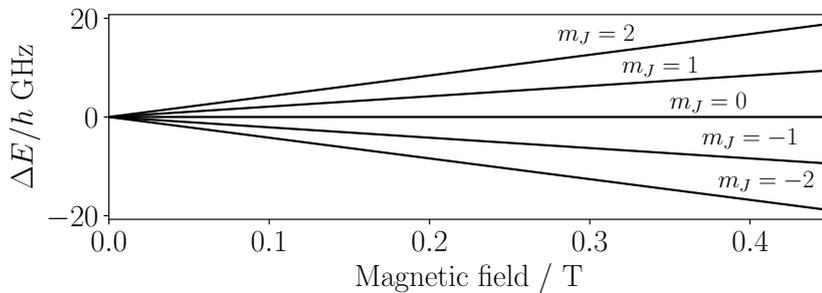


Figure 3.11: Zeeman splitting of  $^3P_2$  argon. The applied field lifts the degeneracy of the five sub-levels labelled by  $m_J$  giving two high-field seeking, two low-field seeking, and one level not affected by the field.

Because the energy changes linearly with respect to the magnetic field the gradient of the energy can be directly related to the gradient of the field, so that the force exerted on a particle is given by:

$$F_{[x,y,z]} = -g_J \mu_B m_J \frac{\partial |\mathbf{B}|}{\partial [x, y, z]}. \quad (3.4.1)$$

For argon, therefore, the  $m_J = 1$  and  $2$  states are low-field seeking and the  $m_J = -1$  and  $-2$  are high-field seeking. This is incorporated into our simulations of the travelling-wave Zeeman decelerator. The simulation code generates a sample of atoms with initial positions, velocities and projection quantum numbers  $m_J$  as described in sections 3.2 and 3.5. The simulation is step-wise, and at each step the particle motion is evaluated. If the decelerator field is inactive the sample positions are incremented according to simple free flight. When the field is active, the magnetic field is calculated using the code described in section 3.6.1; this returns the magnitude  $|\mathbf{B}|$  across a three-dimensional grid.

If the simulation is one-dimensional then the  $z$ -component of the field is extracted and loaded into a spline interpolator (see section 3.7.1) which is queried to return the gradient  $d B_z / dz$  at the location of each particle. For a full three-dimensional simulation we instead load the calculated field into our tricubic interpolator which returns the gradients along all three coordinate axes for each particle in the sample. As described in section 3.7.2 this code library was created specifically to provide a

fast and accurate method of interpolating data across a regular grid, including derivatives, for large query sizes, since a suitable package was not available elsewhere. The tricubic code has been key to much of the work described in this thesis. Earlier simulations of the experiment utilising the first implementation of the tricubic method relied on using a pre-calculated time-averaged decelerator field and scaling the particle positions [2], which was necessary because both setting up and querying the interpolator were too slow to allow on-the-fly calculations of the decelerator field to be used. Once the field gradient is known at the particle positions, the force on each and thus its acceleration is found via equation 2.3.9, and the motion can be solved using a numerical integrator. We originally implemented a fourth-order Runge-Kutta method but subsequently switched to the velocity Verlet algorithm [186], which provides more than adequate accuracy with fewer calculation steps; details are given in section 3.8.

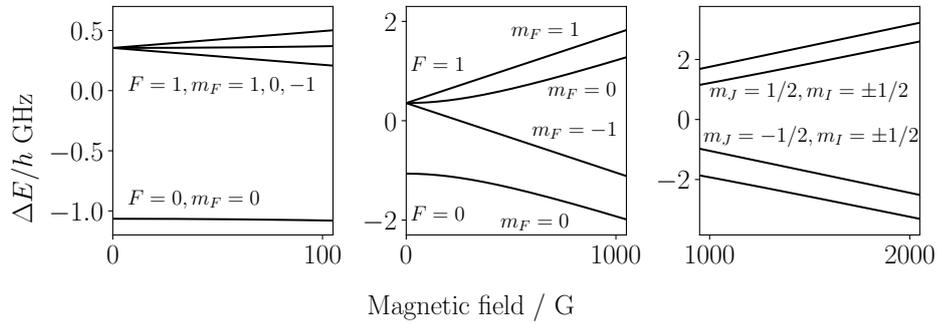


Figure 3.12: Zeeman splitting of ground-state hydrogen, shown for different scales of applied magnetic field  $B$ . Left, at low field a series of energy levels labelled by  $F$  and  $m_F$  and linear in  $B$ . Centre, at intermediate fields not all sub-levels are linear with  $B$ . Right, in the hyperfine Paschen-Back regime the system is described by  $m_I$  and  $m_J$ , and is again linear.

Other species require different approaches, for example figure 3.12 shows the Breit-Rabi diagram for  $^2S_{1/2}$  hydrogen for weak, intermediate and strong fields. In weak fields, in a similar fashion to the coupling of  $\mathbf{L}$  and  $\mathbf{S}$  which gives rise to the total electronic spin  $\mathbf{J}$ , the nuclear and electronic spins  $\mathbf{I}$  and  $\mathbf{J}$  interact to give a total spin  $\mathbf{F} \equiv \mathbf{I} + \mathbf{J}$ .  $m_I$  and  $m_J$  are not good quantum numbers and instead we use  $F$  and  $m_F$ , with  $F = 0, m_F = 0$  or  $F = 1, m_F = -1, 0$  or  $1$ . A common approach [85] is to write the Zeeman Hamiltonian in terms of  $\mathbf{F}$  as:

$$\hat{H}_{Zee} = \frac{\mu_B g_F \hat{F}_z B_z}{\hbar}, \quad (3.4.2)$$

which gives energy shifts of:

$$\Delta E = g_F \mu_B m_F B_z, \quad (3.4.3)$$

giving the force on an atom as dependent on  $m_F$ :

$$F_{[x,y,z]} = -g_F \mu_B m_F \frac{\partial |\mathbf{B}|}{\partial [x, y, z]}. \quad (3.4.4)$$

For fields strong enough that the Zeeman shift is larger than the hyperfine splitting  $IJ$ -coupling breaks down and we enter the hyperfine Paschen-Back regime in which the system can be described in terms of  $m_I$  and  $m_J$ , both taking values  $\pm 1/2$ . This begins at field strengths approximately given by the ratio  $A_{\text{hfs}}/\mu_B$  where  $\mu_B$  is the Bohr magneton and  $A_{\text{hfs}}$  is the hyperfine coupling constant,

which varies by species / state. For hydrogen this equates to roughly 0.1 T. As was the case for metastable argon equation 3.4.1 can be applied to solve hydrogen motion in fields with appropriate magnitudes.

However, modelling hydrogen in the Zeeman decelerator presents problems. If it is only required to consider the maximally spin-stretched  $F = 1$ ,  $m_F = 1$  state then either of equations 3.4.1 or 3.4.4 could be used, since the Stern-Gerlach type interaction only depends on the gradient of the energy shift, but adding in the other sub-levels to get a more complete prediction of the experimental outcome is more difficult. An individual trap in the decelerator field has field magnitudes with a maximum of order  $\approx 0.5$  T but approaching zero at the centre, so a hydrogen atom caught in the trap during the operation of the experiment might experience fields spanning more than one of the regimes shown. Atoms with  $F = 0$  or 1,  $m_F = 0$  moving from the centre to the extrema of the trap would go from experiencing zero force to having high / low-field seeking behaviour. To obtain a more complete solution it is necessary to include the hyperfine interaction when calculating the energy eigenvalues of the different states in the sample. This interaction between the nuclear spin and electronic angular momentum can be well approximated as a multipole expansion [187]; the contribution of the terms in the series tails off rapidly and generally only the first two, the dipole and quadrupole, need to be considered. The quadrupole term also vanishes unless  $I \geq 1$  and  $J \geq 1$ , and may often be neglected if very precise calculations are not required. The approximate Hamiltonian of the dipole moment is:

$$\hat{H}_{\text{Dip}} = \frac{A_{\text{hfs}}}{\hbar^2} \hat{\mathbf{I}} \cdot \hat{\mathbf{J}}. \quad (3.4.5)$$

Combining this with equation 2.3.7 gives the total Hamiltonian  $\hat{H}_{\text{Tot}} = \hat{H}_{\text{Zee}} + \hat{H}_{\text{Dip}}$ :

$$\hat{H}_{\text{Tot}} = \frac{A_{\text{hfs}}}{\hbar^2} \hat{\mathbf{I}} \cdot \hat{\mathbf{J}} + \frac{\mu_B g_J \hat{\mathbf{J}} \cdot \mathbf{B}}{\hbar}. \quad (3.4.6)$$

Since we are working in field regimes such that the energies are much smaller than the fine structure we neglect coarse and fine structure and choose our zero energy to correspond to the zero-field limit of the Zeeman Hamiltonian, and use the Hamiltonian in equation 3.4.6 for the whole system. For  $J = 1/2$  the standard solution is the ‘Breit-Rabi formula’ [187]:

$$\Delta E_{F=I \pm 1/2} = -\frac{\Delta W}{4(I + 1/2)} \pm \frac{\Delta W}{2} \sqrt{\left(1 + \frac{2m_F x}{I + 1/2} + x^2\right)}, \quad (3.4.7)$$

where  $x \equiv \mu_B g_J B_z / \Delta W$  and  $\Delta W$  is the separation between the two hyperfine sub-levels at zero field, related to the hyperfine structure constant by  $A_{\text{hfs}} \equiv \Delta W / (I + 1/2)$ . The latter term arises from the Fermi contact interaction [188]; for hydrogen  $I = 1/2$  and the correction vanishes, but it is significant for other species. The Breit-Rabi formula gives rise to four separate equations for the Zeeman shift of the sub-levels in hydrogen, but in order to calculate trapping forces the gradients of the energy are required. For a  $J = 1/2$  system we can proceed by calculating the Zeeman shift at each point in our grid using the Breit-Rabi equations and then interpolating for the first derivatives, however for  $J \neq 1/2$  a more general approach to solving the Hamiltonian in equation 3.4.6 is required. One method is to work in the coupled basis, beginning with the zero-field eigenfunctions. The required Zeeman Hamiltonian matrix elements need to be calculated in the basis  $|F, m_F\rangle$ . For ground-state hydrogen  $F = 0$  or 1,  $L = 0$  so  $J = S = 1/2$  and the matrix elements have a form

$\langle F'm'_F | \hat{S}_z | F, m_F \rangle$ . The operator for  $S_z$  acts on the  $|S, m_S\rangle$  basis, so  $|F, m_F\rangle$  must be decomposed to the  $|S, I, m_S, m_I\rangle$  basis by calculating the Clebsch-Gordan coefficients.

Working in the coupled basis can become quite complex and cumbersome for some species / states, so instead we assemble a matrix of the Hamiltonian in terms of the uncoupled basis  $|I, J, m_I, m_J\rangle$ , then diagonalise it to find the energy eigenvalues as a function of external field. In order to construct the Hamiltonian where both  $\mathbf{I}$  and  $\mathbf{J} = 1/2$  the Pauli spin matrices  $\boldsymbol{\sigma}$  could be used, with  $\hat{\mathbf{I}} = \hbar\hat{\boldsymbol{\sigma}}_I/2$  and  $\hat{\mathbf{J}} = \hbar\hat{\boldsymbol{\sigma}}_J/2$ . In order to more conveniently deal with larger angular momenta we can instead utilise the ladder operators  $\hat{s}_+ \equiv \hat{s}_x + i\hat{s}_y$  and  $\hat{s}_- \equiv \hat{s}_x - i\hat{s}_y$  which are easily generated on a computer for an arbitrary value of spin, and then used to give the components  $s_x = \hbar/2(s_- + s_+)$ ,  $s_y = i\hbar/2(s_- - s_+)$  and  $s_z = \hbar/2(s_+s_- - s_-s_+)$ . To create the hyperfine Hamiltonian, the ladder operators are used to generate  $i_{[x,y,z]}$  and  $j_{[x,y,z]}$ ; for hydrogen both components are spin 1/2 so these are the standard Pauli matrices. We then create the combined state space combining these two spin spaces by taking their tensor products  $I_{[x,y,z]} = j_{[x,y,z]} \otimes \mathbb{I}_i$  and  $J_{[x,y,z]} = \mathbb{I}_j \otimes j_{[x,y,z]}$  where  $\otimes$  is the Kronecker product and  $\mathbb{I}$  is an identity matrix of dimensionality of the opposite component. The dipole Hamiltonian can then be assembled, in the basis  $|1/2, 1/2\rangle, |1/2, -1/2\rangle, |-1/2, 1/2\rangle, |-1/2, -1/2\rangle$ :

$$\hat{H}_{\text{Dip}} = I_x J_x + I_y J_y + I_z J_z = A_{\text{hfs}} \begin{bmatrix} 1/4 & 0 & 0 & 0 \\ 0 & -1/4 & 1/2 & 0 \\ 0 & 1/2 & -1/4 & 0 \\ 0 & 0 & 0 & 1/4 \end{bmatrix}. \quad (3.4.8)$$

The Zeeman Hamiltonian can be written in terms of  $J_z$ :

$$\hat{H}_{\text{Zee}} = \frac{g_J \mu_B B_z J_z}{\hbar} = g_J \mu_B B_z \begin{bmatrix} 1/2 & 0 & 0 & 0 \\ 0 & -1/2 & 0 & 0 \\ 0 & 0 & 1/2 & 0 \\ 0 & 0 & 0 & -1/2 \end{bmatrix}, \quad (3.4.9)$$

with the Landé factor  $g_J = 1 + (J(J+1) - L(L+1) + S(S+1))/(2J(J+1))$ . It is worth noting that we have here used the convention that the magnetic field  $\mathbf{B}$  points along an arbitrary  $z$ -axis with magnitude  $B_z$ , with the projection of  $\mathbf{I}$  and  $\mathbf{J}$  along the same axis. This is not the same frame of reference as the physical experiment being simulated, in which we choose  $z$  to be the direction of propagation of the molecular beam. At each step when the eigenenergies are calculated the magnitude of  $\mathbf{B}$  is used, rather than its projection on the physical  $z$ -axis.

The dipole and Zeeman matrices are constructed once at the beginning of the simulation, then at each timestep the decelerator field is found at each point of the simulation grid and used to generate the total Hamiltonian  $\hat{H}_{\text{Tot}} = \hat{H}_{\text{Dip}} + \hat{H}_{\text{Zee}}$ . This is then solved using the SciPy [179] ‘linalg’ library. This produces a grid of Zeeman energies for each state present, *e.g.* for hydrogen we would have four fields, each of which can be loaded into the tricubic interpolator. In the previous example we used the interpolator to find the gradients of the decelerator field, whereas now we are able to return the gradient of the energy and thus the force exerted by the field at any point. The Zeeman shift code library has been developed to be easily adapted to other species, and has been used in modelling the behaviour of hydrogen and lithium in a hybrid magneto-optical trap [189].

### 3.5 Source State Population of Argon

Section 3.2 described the simulation of our supersonic expansion source, in which a dielectric-barrier discharge (DBD) is utilised to produce metastable atoms. This can be applied to several different species but in this work we are largely concerned with argon, which was used to characterise the experiment. The DBD is necessary for two reasons; firstly, ground-state argon atoms are not energetic enough to trigger the microchannel-plate detector (MCP), and secondly, we require the atoms to be in a paramagnetic state in order for them to be manipulated by the decelerator fields. Consideration of the quantum state of the atoms was not necessary for the free-flight model but is essential when including the effects of the decelerator fields. Figure 3.13 shows the lowest excited states of argon, of which only the triplets  $^3P_1$  and  $^3P_2$  have sub-levels with non-zero magnetic moments; and of these only  $^3P_2$  is metastable [176].

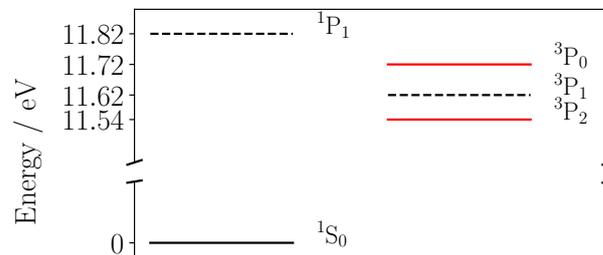


Figure 3.13: Energy levels of argon in the first excited configuration, relative to the ground state. Dashed lines indicate states with a short excited lifetime before undergoing radiative decay, states in red are metastable.

The  $^3P_2$  state has five sublevels, which are degenerate in zero magnetic field. The Breit-Rabi diagram in figure 3.14 shows the splitting of these states in an external field, with the optimal state for magnetic trapping and deceleration shown in red; see section 3.4. When simulating the supersonic expansion source (section 3.2) each atom in the initial distribution is randomly assigned a value of the  $m_J$  quantum number from the set  $[-2, -1, 0, 1, 2]$ . Further discussion of the source population is given in section 3.10.

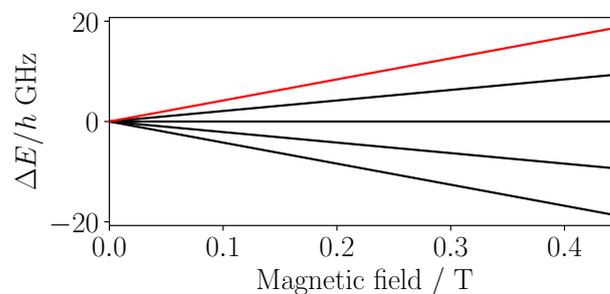


Figure 3.14: Energy splitting of the sublevels of the  $^3P_2$  state of argon in an applied magnetic field. The maximally spin-stretched low-field-seeking  $m_J = 2$  sublevel is shown in red. See also figure 3.11.

## 3.6 Calculating Magnetic Fields

### 3.6.1 Decelerator Fields

In order to simulate the travelling-wave Zeeman decelerator, we need to be able to calculate the magnetic field produced by the deceleration coils at the position of each particle in the beam, at every time step. Section 2.3.6 described how we determine the sequence of voltage pulses necessary to approximate the target current waveform in the decelerator coils, and the same code is the basis of our field calculations.

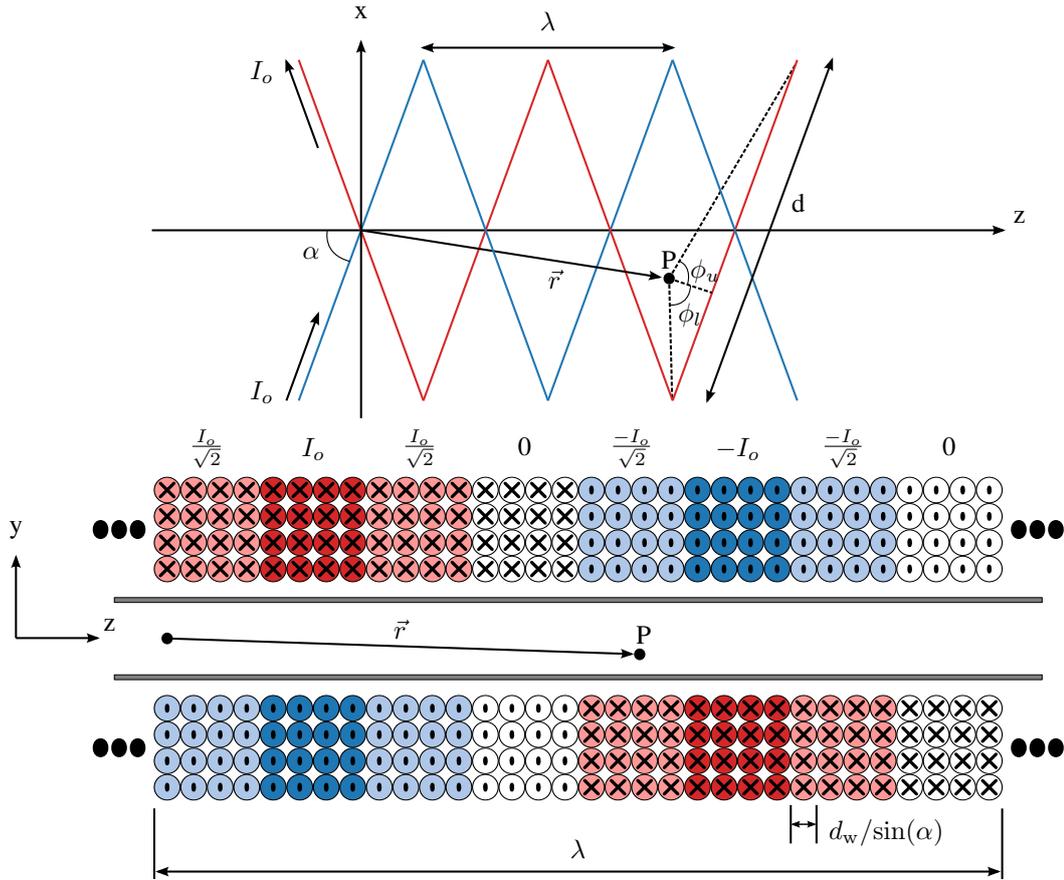


Figure 3.15: Calculating the magnetic field in the decelerator. Upper panel, simplified structure of a coil from above. At a point  $P$  the field due to the nearby wire segment is calculated via the Biot-Savart law. This calculation is carried out across a spatial grid surrounding the region of interest for all wires, then summed together. Lower panel, side-view cross-section through the coils. The direction and magnitude of the current as a fraction of the peak  $I_0$  is shown, where neighbouring phases have a phase offset of  $\pi/4$  radians. *cf.* section 2.3.4. Image reproduced from L. McArd [2].

Figure 2.10 shows an example of the PWM calculation, with the target and actual current over time. This sequence can be used to determine the magnetic field produced by the wire segment in the coil via the Biot-Savart law:

$$\mathbf{B} = \frac{\mu_0}{4\pi} I \int \frac{\mathbf{R} \times d\mathbf{s}}{|\mathbf{R}|^3}, \quad (3.6.1)$$

where  $d\mathbf{s}$  is the current element and  $\mathbf{R}$  is the displacement vector from  $P$  to the wire. In order to

apply this we work in a coordinate system rotated by the angle  $\alpha$ , such that the wire lies along one axis. The direction of the field at  $P$  is given by the right-hand rule for the cross product  $\mathbf{R} \times d\mathbf{s}$  and the magnitude  $B$  can be found from the Biot-Savart law for finite-length conductors [190]:

$$B = \frac{\mu_0}{4\pi d} I \int [\sin(\phi_u) - \sin(\phi_l)], \quad (3.6.2)$$

where  $d$  is the distance between  $P$  and the axis along which the wire is aligned, and the angles  $\phi_u, \phi_l$  are as shown on figure 3.15. The grid of coordinates over which we wish to know the field in the laboratory frame is generated, transformed into the rotated frame, and the corresponding values  $B$  calculated. The direction of the field at each point is also found, rotated back into the laboratory frame, and used to determine the vector components  $\mathbf{B}$ . As seen in the upper row of figure 3.15 each wire segment of a coil can be treated as several individual straight wires; the coils are three periods long meaning each segment can be modelled as six wires.  $\mathbf{B}$  is found for each wire, and then the six fields can be summed together.

The complete field arising from a module can be found by proceeding in the above manner for each wire per phase, each phase per coil, and each coil per module, but we can adopt a more efficient procedure. Each coil is connected as four phases of  $2 \times 4$  wires each, and the coil is paired giving effectively a block of  $4 \times 4$  wires. If at some time  $t_n$  a given wire, say in one of the red blocks shown in 3.15, carries current  $I_n$  then so do all of the wires in that phase. Likewise, all of the wires in the corresponding block in the coils opposite carry  $-I_n$ . Also, equation 3.6.2 shows that  $B \propto I$  so we can calculate the field across our spatial grid due to the current in the same phase of both coil pairs for, say, 1 A, and store this result for later use. We can then determine the grid of values of  $B$  for any arbitrary  $I$  by simply scaling it, rather than calculating the result for each line segment and rotating coordinate systems. Each decelerator section, which contains four separate coils, can be stored as four individual pre-calculated fields. Rather than encompassing the entire beam path we define a grid that covers a region slightly larger than the sample delivery tube in the  $x$ - and  $y$ -directions, and in  $z$  beginning and ending several centimetres either side of the decelerator coils. This reduces the calculation time and size of the output array.

The current in each phase is determined for a range of times as described in section 2.3.6, and a one-dimensional interpolator is used to return the current  $I_n$  for any arbitrary time as required. The number of data points used to describe the time-varying current in the coils is decided in accordance with the Nyquist–Shannon sampling theorem [191]; an analogue function with no frequency component greater than  $F$  can be discretised and completely recovered if sampled at a frequency of no less than  $2F$ .

Earlier simulations of the experiment utilised this method to visualise the decelerator field as a function of time, but were unable to produce detailed simulations of the experiment due to the lack of an efficient method of interpolating the field. Instead, the average field profile of a single magnetic trap from the decelerator was calculated and its movement along the beam axis was modelled via coordinate shifts [2]. This was successful at reproducing the broad structure of our experimental data, but was unable to capture the effects arising from the fluctuations in trap depth and position caused by the PWM process. The improved tricubic interpolator allowed much more detailed dynamic simulations of the complete experiment to be performed, as will be shown in following sections.

### 3.6.2 Static Fields

Section 3.6.1 describes modelling the magnetic fields in the Zeeman decelerator, which we are able to do via the Biot-Savart law due to the coil geometry. Permanent magnets such as those in the decelerator quadrupole or the proposed hybrid magnetic trap / magneto-optical trap (MT-MOT, see chapter 5) have no simple analytic solution, however, and must be modelled differently. This section describes the results we have obtained from using two alternative software packages.

The MT-MOT fields were first calculated [2] using the Radia package [171], which is a plugin developed by the European Synchrotron Radiation Facility for the Mathematica program. Radia is used for magnetostatics modelling, for assemblies containing both permanent and electromagnetic components. Most packages of this type utilise finite element analysis, whereas Radia utilises boundary integral methods. Consider the magnetic field  $\mathbf{H}$  produced by an object of uniform magnetization  $\mathbf{M}$ , this can be represented by the matrix  $\mathbf{Q}$ :

$$\mathbf{H} = \mathbf{QM} \quad (3.6.3)$$

For a line parallel to vector  $\mathbf{v} = (v_x, v_y, v_z)$  and passing through a point  $r_0$  the object produces a field given by the integral:

$$\mathbf{I} \equiv \int_{-\infty}^{\infty} \mathbf{H}(\mathbf{r}_0 + \mathbf{v}_s) ds = \mathbf{GM}. \quad (3.6.4)$$

Equations 3.6.3 and 3.6.4 are quite general, irrespective of the shape of the object. For the example of a rectangular parallelepiped aligned along the Cartesian XY, XZ and YZ planes the components of  $\mathbf{Q}$  are given by:

$$\begin{aligned} Q_{xx} &= \frac{1}{4\pi} \sum_{i,j,k=1}^2 (-1)^{i+j+k+1} \tan^{-1}[x_i^{-1} y_j z_k (v_{i,j,k})^{-1/2}], \\ Q_{xy} &= \frac{1}{4\pi} \ln \left[ \prod_{i,j,k=1}^2 [z_k + (x_i^2 + y_j^2 + z_k^2)^{1/2}]^{-1^{i+j+k}} \right], \\ v_{i,j,k}^2 &= x_i^2 + y_j^2 + z_k^2, \\ x_{1,2} &= x_c - x_0 \mp w_x/2, \\ y_{1,2} &= y_c - y_0 \mp w_y/2, \\ z_{1,2} &= z_c - z_0 \mp w_z/2, \end{aligned} \quad (3.6.5)$$

where  $w_x, w_y, w_z$  are the dimensions and  $x_c, y_c, z_c$  the coordinates of the centre of the block,  $x_0, y_0, z_0$  is the reference point of interest. The other components of  $\mathbf{Q}$  are found by interchanging coordinates in equations 3.6.5. A similar set of equations apply for the matrix  $\mathbf{G}$  for a parallelepiped, for example to integrate along  $v_x = v_y = 0$ :

$$\begin{aligned}
G_{xx} &= \frac{w_y}{2\pi} \sum_{i,k=1}^2 (-1)^{i+k} \tan^{-1}(x_i z_k^{-1}), \\
G_{zz} &= \frac{w_y}{2\pi} \sum_{i,k=1}^2 (-1)^{i+k} \tan^{-1}(x_i^{-1} z_k), \\
G_{xz} &= \frac{w_y}{4\pi} \ln \left[ \frac{(x_1^2 + z_2^2)(x_2^2 + z_1^2)}{(x_1^2 + z_1^2)(x_2^2 + z_2^2)} \right], \\
G_{yy} &= G_{xy} = G_{yz} = 0.
\end{aligned} \tag{3.6.6}$$

Other volume shapes have similar expressions. Completely arbitrary shapes are built up from many smaller volumes of known systems. Radia contains similar objects for modelling current-carrying conductors. For a complete system a total interaction matrix is created and solved as required to provide the necessary field components. Our group previously sampled the field around a mock-up of the MT-MOT using permanent magnets with a Hall probe and compared the results with a Radia model, finding good agreement. Other work by the author used the Radia model of the MT-MOT in Monte Carlo simulations, in work assessing the feasibility of loading the trap with a dual-species beam of lithium and calcium monohydride [151]. It was found that Radia introduces erroneous artefacts at the boundaries of curved permanent magnets, and care had to be taken to manually check for and remove these features to avoid introducing errors into the simulations.

An alternative method of calculating the fields around the toroidal magnets of the MT-MOT was developed using the FEMM package (Finite Element Method Magnetics) [172]. This supports modelling systems comprising any combination of steady current-carrying and permanent magnetic elements, with the caveat that it is two-dimensional. FEMM handles magnetostatic (time-invariant) problems in the following way; the flux density  $\mathbf{B}$  and field intensity  $\mathbf{H}$  must satisfy:

$$\nabla \cdot \mathbf{B} = 0, \tag{3.6.7}$$

and:

$$\nabla \times \mathbf{H} = \mathbf{J}, \tag{3.6.8}$$

where  $\mathbf{J}$  is the current density, with  $\mathbf{B}$  and  $\mathbf{H}$  related by some permeability  $\mu$ :

$$\mathbf{B} = \mu \mathbf{H}. \tag{3.6.9}$$

$\mathbf{B}$  can be written in terms of the magnetic vector potential  $\mathbf{A}$  in a form that satisfies equation 3.6.7:

$$\mathbf{B} = \nabla \times \mathbf{A}, \tag{3.6.10}$$

giving:

$$\nabla \times \left( \frac{1}{\mu} \nabla \times \mathbf{A} \right) = \mathbf{J}. \tag{3.6.11}$$

For non-linear materials  $\mu$  will be some function of  $B$ , and can be replaced in equation 3.6.11 by the appropriate relationship. A magnetic element is created in FEMM by specifying shape, dimensions, material composition and magnetisation vector. The software uses finite element analysis to find solutions to the required potential; the problem domain is broken down into a large number of small, simply-shaped domains in which the ‘true’ solution can be individually closely approximated. In the case of FEMM the regions are triangular, and each element contains a solution found by linear interpolation of the values at the three vertices. The mesh granularity can be increased in order to give more accurate results at the cost of greater computation time. The model also requires boundary conditions, with a variety of configurable options such as Dirichlet or Neumann available. It was found that using the default setting of automatically choosing appropriate boundary conditions was adequate for our permanent magnet calculations. The resultant solution can then be queried using the Lua scripting language.

FEMM is two-dimensional, supporting either a standard  $x, y$  Cartesian grid or an axisymmetric cylindrical polar grid  $r, z$ . To model the MT-MOT the latter was used, with one of the toroidal magnets created in cross-section with the rotation axis passing through its centre. An appropriate mesh size is chosen and FEMM calculates the approximate solution for the magnetic field  $\mathbf{B}$ . This can be queried at any point  $r, z$  to return either the magnitude  $B$  or the vector components  $B_r, B_z$ . The model and its solution are shown in figures 3.16 and 3.17. The material chosen for the magnet is N27 neodymium, with the direction of the magnetisation vector indicated by the arrow; the other region is air.

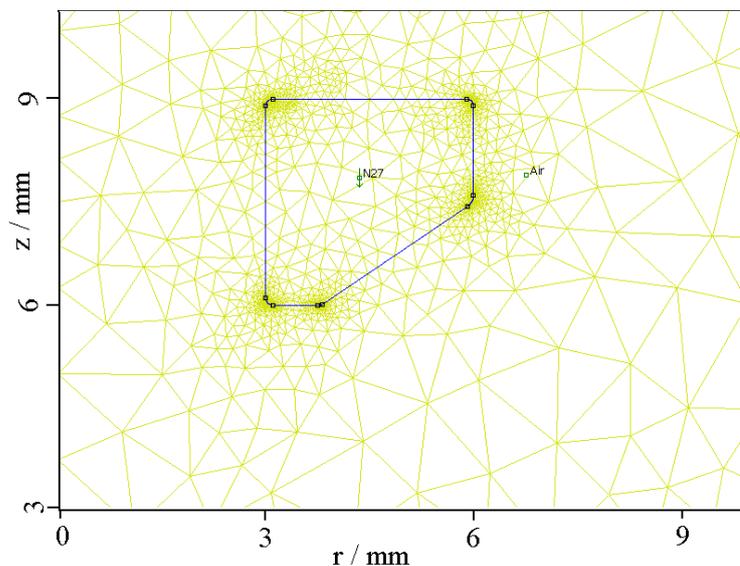


Figure 3.16: Cross-section model of a toroidal ring magnet in FEMM, rotation axis vertical, radius horizontal. The dimensions and materials are specified, including magnetisation vectors - N27 is a grade of neodymium material. The triangular discretization mesh over which the approximate solution is calculated is shown; the mesh is finer where boundaries are more complex.

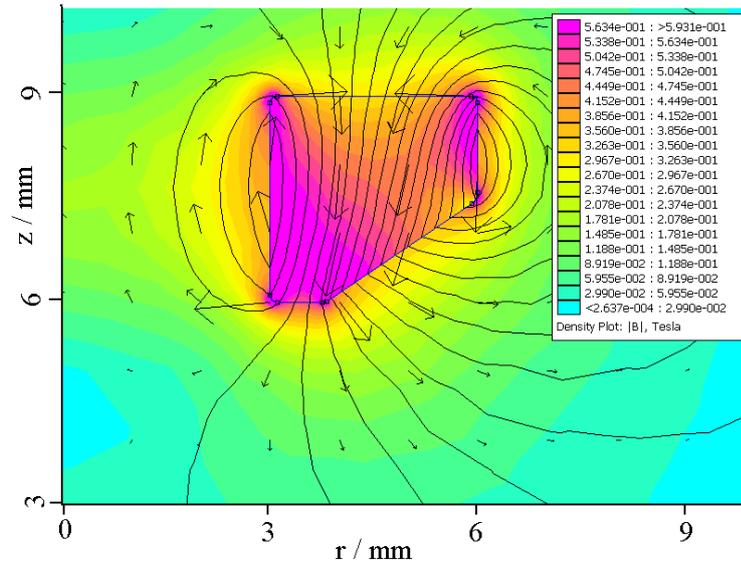


Figure 3.17: Cross-section model of a toroidal ring magnet in FEMM, rotation axis vertical, radius horizontal. Shown is the calculated solution of the model in figure 3.16; flux lines are shown, with the heatmap indicating the field magnitude.

We have developed a method to expand this solution to three dimensions, by exploiting the rotational symmetry of the system. Firstly, we assemble a grid of coordinates  $x, y, z$  across which we wish to know the field. The origin is at the centre of the trap and the modelling axes are already defined, for example  $-50 \text{ mm} \leq z \leq 50 \text{ mm}$  and  $0 \leq r \leq 50 \text{ mm}$ . Choosing the other coordinates as  $0 \leq x, y \leq 50 \text{ mm}$  we can then query the field at each point  $x, y, z$  by calculating  $r = \sqrt{x^2 + y^2}$  (since the coordinate system is cylindrical), returning  $B_r, B_z$  and  $B$ . An interface written in Lua extracts the data as the matrix  $x, y, z, B_r, B_z, B$ , and Python code is used to complete the field. At each point the angle between  $r$  and the  $x$ -axis is given by  $\phi = \arccos(y/x)$  allowing us to calculate  $B_x = B_r \cos(\phi)$  and  $B_y = B_r \sin(\phi)$ . The field components  $B_{[x,y,z]}$  are now known across the grid  $x, y, z$ , for positive values of  $x$  and  $y$ ; the set is mirrored in the  $xz$ -plane and concatenated, and then again in the  $yz$ -plane. This gives the vector field produced by a single toroidal magnet across the entire region of interest. A copy of the data set is then rotated by  $\pi$  about  $y$  to give the opposite magnet with the resultant field produced by summing them together, giving the pair of magnets aligned along  $z$ . Similar rotations by  $\pi/2$  about the  $y$  and  $x$  axes give the pairs of magnets along  $x$  and  $y$ , and finally the entire set of six magnets can be assembled. This method is simpler to adjust and produces smoother outputs than Radia for rotationally symmetric objects.

## 3.7 Interpolation

Where some quantity of interest is known as a discrete data set, it is necessary to interpolate the set to acquire knowledge of the value at arbitrary points within the data domain. Broadly speaking interpolation schemes fall into two categories; global, in which some approximating function is fitted across the entire data set, and local, in which individual functions are created piecewise between each pair of neighbouring data points. Global schemes such as polynomial interpolation may produce very smooth results but can be computationally expensive to calculate, and often suffer from large oscillations at the boundaries ('Runge's phenomenon' [192]).

The simplest local method is linear interpolation which fits a straight line between each pair of adjacent data points, but this does not produce a very smooth output - better are splines [193], which use low-degree (typically cubic) polynomial functions, defined piecewise between pairs of points. These do not suffer from Runge's phenomenon, are relatively cheap to calculate, are smoother than linear methods and, being local, large errors (often found at data boundaries) do not propagate. Cubic splines have been heavily utilised in the work described in this thesis; there will now follow a brief overview of background theory and a discussion of the algorithms employed (including a novel extension to a well-known method).

### 3.7.1 Spline Interpolation

A 'spline' is one of a large family of related mathematical techniques employing piecewise polynomials. They take their name from thin flexible rulers used in ship and aircraft design; when fixed at key points or 'knots' on a blueprint, these devices naturally adopt a curve that minimises strain along their length. Mathematical splines appeared in the 1940s [194] and were further developed in the automotive industry in the 1950s and 1960s. This was the era of early computer-aided manufacturing, and splines allowed the use of arbitrary curves rather than simple conic sections [195].

Consider data points along an axis  $x$  in the interval  $[x_0, x_n]$ , giving  $n$  sub-intervals:  $[x_0, x_1], [x_1, x_2] \dots [x_{n-1}, x_n]$ . Rather than a single function  $f(x)$  covering the entire domain, each sub-interval has its own function:

$$f(x) = \begin{cases} f_0(x), & x_0 \leq x < x_1 \\ f_1(x), & x_1 \leq x < x_2 \\ \dots \\ f_{n-1}(x), & x_{n-1} \leq x \leq x_n. \end{cases} \quad (3.7.1)$$

The spline is of degree  $k$  if the polynomials have maximum degree  $k$ , *i.e.* for the cubic case here each sub-function has at most a term in  $x^3$ . Piecewise functions have a 'smoothness' defined as the continuity of their derivatives at the junction between intervals; if at point  $x_i$  both  $f_{i-1}(x)$  and  $f_i(x)$  share common derivatives from order zero to  $y$  then there is smoothness  $C^y$ .  $f(x)$  is fully continuous within each sub-interval; explicitly,  $f_i(x)$  is  $C^\infty$  in  $[x_i, x_{i+1}]$ .

A cubic spline has an interpolating function of the form:

$$f_i(x) = a_i + b_i(x - x_i) + c_i(x - x_i)^2 + d_i(x - x_i)^3, \quad (3.7.2)$$

For a given data set the coefficients  $a_i$ ,  $b_i$ ,  $c_i$  and  $d_i$  are calculated for each sub-interval; once these are known it is simple to calculate arbitrary values of  $f_i(x)$  as required [196]. Cubic splines of this type for one- and two-dimensional problems are readily available as software libraries for Python [179], C++, Matlab *etc.*, and are generally well optimised for computational efficiency.

### 3.7.2 Tricubic Interpolation

For systems with more than one variable there exist multivariate schemes that combine interpolation of the data set across these variables. Originally motivated by studies of ocean dynamics [197] Lekien and Marsden describe their ‘tricubic’ technique [198] for either time-dependent two-dimensional flows or three-dimensional time-independent flows. Earlier three-dimensional cubic spline methods treated the problem as three one-dimensional problems [199], whereas the Lekien-Marsden solution efficiently combines them. The interpolated function has  $C^1$  continuity and its partial derivatives are analytically accessible.

The tricubic method assumes a data set known at a series of regular grid points in three dimensions  $x$ ,  $y$  and  $z$ . As long as the data are regular in all three dimensions they do not have to be the same length, but without loss of generality we shall here consider a coordinate mesh that is a cube. This mesh is composed of elements that are also cubes, for each of which the interpolant field  $f$  is known at the eight vertices  $p_1 \cdots p_8$  (this is, of course, for a scalar field - for a vector field the same method is simply applied to each component separately). Inside each element  $f$  is in the form of the cubic:

$$f(x, y, z) = \sum_{i,j,k=0}^3 \alpha_{ijk} x^i y^j z^k. \quad (3.7.3)$$

The coefficients  $\alpha_{ijk}$  must be calculated, and can then be used to return  $f$  for an arbitrary point  $(x, y, z)$  within the element. In order to achieve  $C^1$  continuity across the whole domain, the values of  $f$  and its first derivatives  $\partial f/\partial x$ ,  $\partial f/\partial y$  and  $\partial f/\partial z$  must be continuous across the faces and vertices of the elements, giving 32 constraints for each. The list of coefficients in equation 3.7.3 is of length 64, so there must be an additional 32 constraints; it can be shown [198] that the only valid choice is the set:

$$\left[ f, \frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}, \frac{\partial f}{\partial z}, \frac{\partial^2 f}{\partial x \partial y}, \frac{\partial^2 f}{\partial x \partial z}, \frac{\partial^2 f}{\partial y \partial z}, \frac{\partial^3 f}{\partial x \partial y \partial z} \right] \quad (3.7.4)$$

To see why derivatives such as  $\partial^2 f/\partial x^2$  cannot be used, consider the value of  $\partial^2 f/\partial x^2$  at point  $p_1$ , which we can take to be at the origin. Along the  $x$ -axis equation 3.7.3 reduces to the cubic spline:

$$f(x) = \sum_{i=0}^3 \alpha_i x^i. \quad (3.7.5)$$

The values of  $f$  and  $\partial f/\partial x$  at points  $p_1 = (0, 0, 0)$  and  $p_2 = (1, 0, 0)$  are fixed, forming a unique spline. This also constrains  $\partial^2/\partial x^2$  at  $p_1$  and  $p_2$ , which is therefore not independent of the values of  $f$  and  $\partial f/\partial x$ . More generally this applies to all derivatives of the form  $\partial^2/\partial z^2$ ,  $\partial^3/\partial y \partial z^2$ , *etc.*, and means that no cubic spline method can achieve  $C^2$  continuity [200]. In order to obtain higher-order

derivatives the approximating function must also be of higher order, such as the  $C^2$  ‘triquintic’ spline recently described by Bradach [201].

The derivatives in set 3.7.4 are found via finite-difference methods [202]. If the coefficients  $\alpha_{ijk}$  from equation 3.7.3 are placed into a vector  $\alpha$  and the values of  $f$  and its derivatives from set 3.7.4 into another vector  $\mathbf{b}$ , they are related by a  $64 \times 64$  matrix  $B$ :

$$B\alpha = \mathbf{b}, \quad (3.7.6)$$

where the components of  $B$  are integers. If the elements are normalised to be unit cubes during calculations, this matrix is the same for all elements, and so only needs to be calculated once. During interpolation the results are scaled back to their actual values. The matrix is invertible, so:

$$B^{-1}\mathbf{b} = \alpha, \quad (3.7.7)$$

allowing the coefficients  $\alpha$  to be calculated. In order to perform an interpolation for an arbitrary point  $(x, y, z)$ , firstly the appropriate volume element containing the query point is located. The alpha coefficients for this element are calculated, and then combined with the query coordinates in equation 3.7.3 to return  $f$ . The coefficients  $\alpha_{ijk}$  are a tensor of order 3, which is applied to the vectors  $(1, x, x^2, x^3)$ ,  $(1, y, y^2, y^3)$  and  $(1, z, z^2, z^3)$ . Differentiating these vectors can be used to return the derivatives of  $f$  - for example, using the set  $(0, 1, 2x, 3x^2)$ ,  $(1, y, y^2, y^3)$  and  $(1, z, z^2, z^3)$  gives  $\partial f / \partial x$ .

### 3.7.3 Quadcubic Interpolation

We have extended the Lekien Marsden method to four dimensions, allowing a time-dependent three-dimensional field (or time-independent four-dimensional field) to be interpolated directly. Instead of a cube in  $(x, y, z)$  with eight vertices, each volume element in our interpolation region is a tesseract in  $(x, y, z, t)$  with 16 vertices, inside of which we have the function:

$$f(x, y, z, t) = \sum_{i,j,k,l=0}^3 \alpha_{ijkl} x^i y^j z^k t^l. \quad (3.7.8)$$

There are 256 coefficients  $\alpha_{ijkl}$  meaning 256 constraints are required, and the derivatives given in set 3.7.4 only supply 128 of them. Again noting the requirement to add independent constraints we add the additional derivatives in set 3.7.9 to set 3.7.4:

$$\left[ \frac{\partial f}{\partial t}, \frac{\partial^2 f}{\partial x \partial t}, \frac{\partial^2 f}{\partial y \partial t}, \frac{\partial^2 f}{\partial z \partial t}, \frac{\partial^3 f}{\partial x \partial y \partial t}, \frac{\partial^3 f}{\partial x \partial z \partial t}, \frac{\partial^3 f}{\partial y \partial z \partial t}, \frac{\partial^4 f}{\partial x \partial y \partial z \partial t} \right] \quad (3.7.9)$$

In order to calculate the derivatives using finite-differences, the values of  $f$  at the vertices of neighbouring elements is needed - a point in this space is surrounded by  $3^4 = 81$  tesseracts defined by 256 points (since many of the vertices are shared). Placing these in a vector  $\mathbf{x} = c_1 \cdots c_{256}$ , the finite-difference matrix  $D$  returns the elements of set 3.7.4 as a vector  $\mathbf{b}$ :

$$D\mathbf{x} = \mathbf{b}. \quad (3.7.10)$$

The interpolation matrix  $B$  is  $256 \times 256$  in size. For the three-dimensional case Faust *et al.* [203]

describe a reduction in computation time which we have extended to four dimensions; we combine the matrices  $D$  and  $B^{-1}$  to produce a new matrix  $DB^{-1} = A$ , which only needs to be calculated once and then reused as required:

$$A\mathbf{x} = \boldsymbol{\alpha}. \quad (3.7.11)$$

$\alpha_{ijkl}$  is a tensor of order 4; applied to the four-dimensional vectors  $(1, x, x^2, x^3)$ ,  $(1, y, y^2, y^3)$ ,  $(1, z, z^2, z^3)$  and  $(1, t, t^2, t^3)$  it returns  $f$ . As before, taking the derivatives of one of these vectors with respect to its variable allows the interpolated values of  $\partial f/\partial x$ ,  $\partial f/\partial t$  etc. to be calculated.

Simulations of the travelling-wave Zeeman decelerator experiment utilising the quadcubic method were developed. Due to the large number of time steps over which the decelerator field needed to be calculated, however, the codes were found to require an impracticable amount of computer memory; the simulations of the experiment described in chapters 3 and 4 instead utilised the tricubic interpolation scheme. Simulations of the proposed MT-MOT in chapter 4 and in Booshanam [189] also used the tricubic code, since the trap field was not time-dependent. The quadcubic scheme has been applied to simulations of the MT-MOT in which the field has extra time-dependent elements, in an extension of work reported in Walker [151]; these results are not included in this thesis.

### 3.7.4 Software Implementation

There are many choices of programming language available for scientific computing such as Fortran, C / C++, Java, Matlab *etc.* but for some years now one of the most prevalent has been Python [204]. Python is object-oriented, highly compatible, and easy to learn, and being open source means that it is free to use and any available code library can be inspected and modified as needed. Being an interpreted rather than a compiled language means that it can be comparatively slow compared to other options, but this is mitigated by the existence of numerical libraries of ‘wrapped’ fast C or Fortran code which can be natively accessed by a Python script. Examples are SciPy, which contains various fundamental scientific and mathematical features [179] and NumPy, an N-dimensional array library [181].

In order to simulate our travelling-wave Zeeman decelerator and hybrid magnetic trap / magneto-optical trap we needed to solve the motion of a paramagnetic particle in a non-uniform and time-dependent magnetic field. An analytic solution for the field was available in neither case and so the field was calculated on a three-dimensional grid, which we needed to interpolate. One- and two-dimensional interpolators are included in SciPy but no suitable three-dimensional code was available. A version of the Lekien-Marsden tricubic spline [198] was created in Python by U. Krohn for McArd [2], allowing us to interpolate for the magnitude and gradient of a three-dimensional magnetic field. (Although the following discussion of the software describes its use for this specific application, it can be used for any suitably-formatted input).

The two major parts of the tricubic method are, firstly; for a given magnetic field  $B$  determining the matrix of interpolation coefficients  $\alpha$ . For a model of the magnetic trap / magneto-optical trap measuring  $20 \text{ mm} \times 20 \text{ mm} \times 30 \text{ mm}$  with a  $B$ -field determined on a regular grid with  $0.5 \text{ mm}$  spacing the matrix would take up to an hour to calculate on a modern desktop PC. For the decelerator the region of interest around several coil modules is several times larger, with corresponding increase in

computation times, and additionally, every time the field changes the process needs to be repeated. For a time-dependent simulation of particles in a static magnetic field the method worked well but in order to model the molecular beam in the decelerator field a time-averaged potential had to be used [2]. (The interpolation matrix could be saved to a file for later use but this proved prohibitively expensive in total calculation time and disk storage).

The second major part of the interpolation scheme is determining the interpolated values of  $B$  and its derivatives  $\partial B/\partial x$ ,  $\partial B/\partial y$  and  $\partial B/\partial z$  for a given set of coordinates. The overall time taken to do this increases drastically as more particles are included in the simulation, since the first version of the software returned the results for one coordinate at a time. All simulations shown in this thesis use a minimum of  $10^6$  particles, which was impractical with the original method.

The tricubic code was completely rewritten by the author using Numpy [181], a library of array manipulation routines based on optimised Fortran code. Each individual operation was analysed separately to identify bottlenecks in computation, memory access or disk input / output. Firstly, the query code was altered to allow multiple points to be sampled at once; using NumPy for matrix operations rather than Python list comprehension led to an increase in code execution speed of almost two orders of magnitude.

Adopting a standard normalised coordinate system within each volume element of the interpolation region allowed the same coefficient matrix to be used; additionally this could be combined with the finite-differences matrix used to approximate the field derivatives needed to form the function constraints [198, 203]. Rather than calculate all interpolation coefficients in advance the code only starts populating them as they are required and re-uses them where possible, thus saving much redundant operation. The code methods used to determine which volume element a query point lay within were also optimised for block operation. With the rewritten code it was now feasible to run the interpolator in an ‘on-the-fly’ way; a new magnetic field geometry could be fed to the code and it would be able to return queries from it within seconds, allowing us to model the operation of the travelling-wave Zeeman decelerator much more accurately than before.

Additionally, the first version of the code was only written to return the interpolated gradient of a scalar field, since that is the only relevant factor in magnetic trapping. The new version was created so that it could be supplied with either a scalar field, in which case it could return the interpolated values and gradients, or a vector field, in which case it could also return the values and gradients of the magnitude, or the interpolated vector components, or both. When, for example, modelling laser-cooling with magnetic trapping or Majorana transitions, all of these quantities must be known, and the code library natively produces them.

The software has been designed to be modular and easily importable into other projects and is available online under a GPL open-source licence [205, 206]. Both the tricubic and quadcubic codes have been adopted for use in a diverse range of applications, being cited in work in fields such as animal vision [207], quantum computation [208], cosmology [209] and chemistry [210].

### 3.7.5 Accuracy Of The Tricubic Method

The accuracy of the interpolator is difficult to quantify, because it strongly depends on the nature of the data being interpolated. In this section we will consider two simple example cases with analytic solutions, allowing us to calculate estimates for the error in the method. Some caveats regarding the application of the interpolator to different systems will be shown.

The quadrupole electric potential  $V$  due to four alternating charges  $q$  arrayed in a square is given by [211]:

$$V(\mathbf{r}) = \frac{1}{4\pi\epsilon_0|\mathbf{r}|^3} \sum_{i,j} \frac{1}{2} q_{ij} \hat{r}_i \hat{r}_j. \quad (3.7.12)$$

Figure 3.18 shows a plot of a plane through such a potential. The left and centre panels are generated by sampling the function in equation 3.7.12 across uniform grids of, respectively,  $121 \times 121$  and  $13 \times 13$  points. The central sparse grid is then interpolated using the tricubic method across the same  $121 \times 121$  grid, to produce the right panel. To the eye, the interpolator is able to successfully reconstruct the missing detail from the sparse data set; to quantify this by we can compare the differences between a range of interpolated and ‘true’ analytic values.

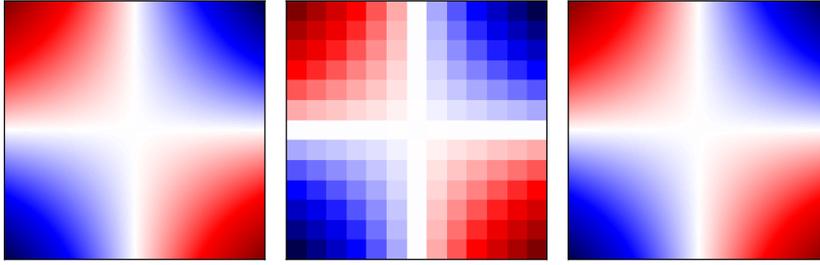


Figure 3.18: The quadrupole electric potential around a series of alternating point charges. Left, the solution is calculated analytically across a square grid with 121 points a side, centre, same as left but 13 points. Right, the sparse grid from the centre is queried with the tricubic interpolator to reproduce the left plot. For this simple model the method is able to reconstruct sparse data well.

We begin by choosing a uniformly-spaced range of  $n$  points  $(x_1, y_1), (x_2, y_2) \dots (x_n, y_n)$  along a diagonal line through the centre of the potential, so that  $(x_1, y_1) = -(x_n, y_n)$ . If the values of the field according to the analytic / interpolated solutions are  $V_A / V_I$  then the root mean square error (RMS) is:

$$E_{\text{RMS}} = \left( \sum_{i=1}^n (V_{A,i} - V_{I,i})^2 / n \right)^{1/2}. \quad (3.7.13)$$

In order to make the error dimensionless we can determine the normalised RMS error by dividing  $E_{\text{RMS}}$  by the mean value  $\bar{V}_{A,i}$ , this can be presented as a fraction or a percentage. (We can work with the mean value even though we are passing through the centre of the field because the potential has the same sign in opposite corners, so the mean is not zero. For a line such as  $x = a$  with  $a \neq 0$  we would have  $\bar{V}_{A,i} = 0$ , so a different approach would be necessary).

The normalised error is calculated for a range of different values of grid spacings along the horizontal and vertical axes, and the value of the error as a function of the spacing can be plotted, as shown in figure 3.19. The grid spacing has also been made dimensionless by presenting it as a percentage of the range  $x_n - x_1$ . As the grid spacing increases so too does the error, although for this simple example it remains small in absolute terms.

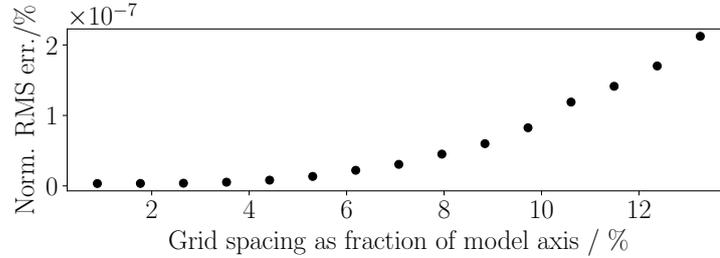


Figure 3.19: Normalised RMS error for tricubic interpolator compared to analytic solution of an electric quadrupole model, as a function of grid spacing. The interpolator is able to reproduce the data accurately with low errors for this simple model.

Magnetic trapping experiments often use a quadrupole field, which can be generated by a pair of coils in an anti-Helmholtz arrangement [46]. The coils are separated by a distance approximately equal to their diameter, and carry equal electrical current but in opposite directions.

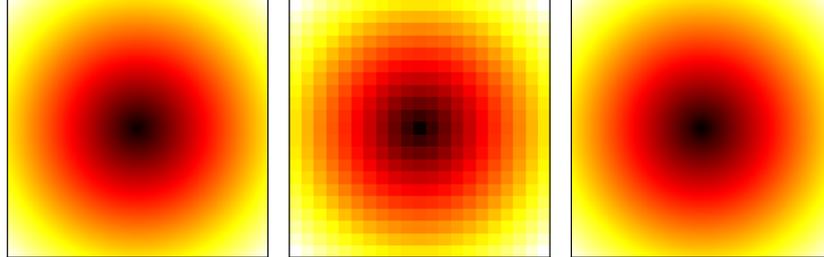


Figure 3.20: The quadrupole magnetic field magnitude in a plane through the centre of a pair of anti-Helmholtz coils. Left, the solution is calculated analytically across a square grid with 201 points a side, centre, same as left but 21 points. Right, the sparse grid from the centre is queried with the tricubic interpolator to reproduce the left plot. Although the method appears to have successfully reconstructed the data, a consideration of the errors reveals problems in the fit at the centre.

In the region around the centre of the trap the magnetic field has a roughly uniform gradient in any direction, falling to zero at the centre. The components of the field  $\mathbf{B}$  can be approximated by:

$$B_{[x,y,z]} = B_0 \begin{bmatrix} x \\ y \\ -2z \end{bmatrix}, \quad (3.7.14)$$

where  $B_0$  is some gradient determined by the current in the coils, and their geometry. Figure 3.20 shows a plot of the field magnitude  $B = |\mathbf{B}|$  in the  $xy$ -plane through such a trap, at  $z = 0$ . As for the previous example, the left and centre panels are generated for fine and coarse grids, in this case of  $201 \times 201$  and  $21 \times 21$  points respectively, and the right plot is produced by interpolating the

centre grid. Again, by inspection it would seem that the tricubic method gives a good reconstruction of the missing detail in the more sparse data set.

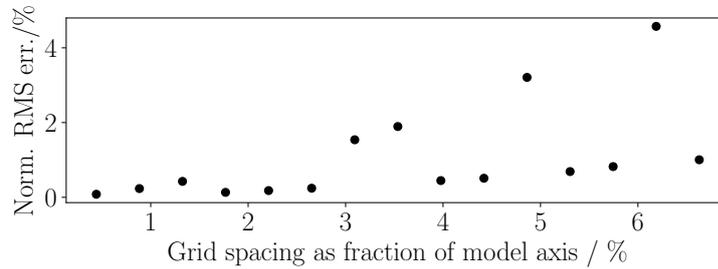


Figure 3.21: Normalised RMS error for tricubic interpolator compared to analytic solution of a magnetic quadrupole model, as a function of grid spacing. Although small overall, the errors fluctuate for certain grid sizes, indicating that the interpolator is not adequately following some feature of the source model at different scales.

We now calculate the error by the same procedure as for the electric quadrupole case; taking points from a line through the centre of the field and comparing the interpolated results to those given by equation 3.7.14. Figure 3.21 shows the normalised RMS error as a function of grid spacing. Although there is the expected general upward trend with increased grid spacing, the values of the errors oscillate, and are orders of magnitude larger than the previous example. The explanation for this is apparent when we compare the outputs of the analytic and interpolated models for particular values of the grid spacing, shown in figure 3.22.

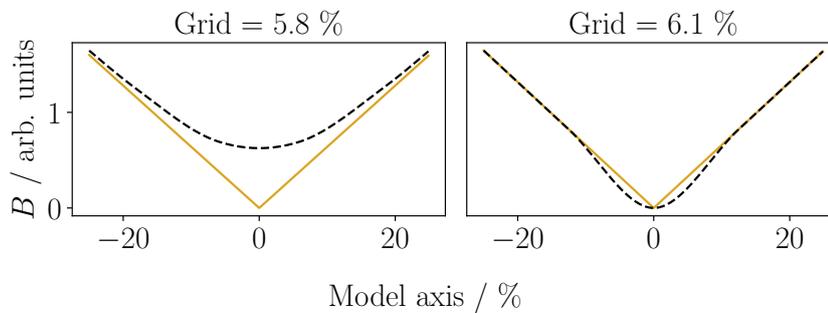


Figure 3.22: The field magnitude  $B$  in the central region of the trap for two values of the interpolation grid spacing. Orange, analytic solution, dashed line, interpolated. The disagreement arises because the tricubic method favours smoothness over accuracy. This explains the oscillations in figure 3.21.

As stated in section 3.7.2, the tricubic method favours smoothness of the interpolated function, and fails to accommodate the sharp change in the field magnitude. As the grid spacing is adjusted the fit alternates between the behaviours seen in the left and right panels of figure 3.22. Rather than working with the field magnitude  $B$  we can avoid this problem by using the vector components  $B_{[x,y,z]}$ , which vary smoothly through the origin. Our code can work with vector fields by separately interpolating the components (see section 3.7.4), and we can then convert to magnitude. The result is shown in figure 3.23; the differences between the interpolated and analytic solutions are now negligible, with the remaining variation arising from floating-point limitations. The two examples given in this section emphasise the point that care must be taken when applying an interpolator that it is an appropriate method for the input data, and that problems may not be apparent at first sight.

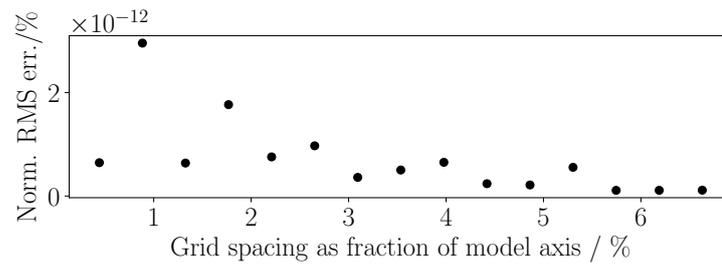


Figure 3.23: As per figure 3.21, but interpolating for the field components rather than magnitude. The error compared to the analytic solution is negligible, with the oscillations arising due to floating-point limitations.

## 3.8 Numerical Integration

In section 3.2 we described a step-wise Monte Carlo simulation approach to reproducing the signal arising from a free-flight supersonic expansion. With no forces involved the displacement of each particle from its position  $\mathbf{r}_n$  at time  $t_n$  is given by  $\mathbf{r}_{n+1} = \mathbf{r}_n + \mathbf{v}\delta_t$ , where  $\mathbf{v}$  is its velocity and the timestep  $\delta_t = t_{n+1} - t_n$ . The free flight model could be reproduced without using a timestep-based simulation, but doing so allows interactions with external forces to be included, since such forces will generally need to be solved using numerical integration.

### 3.8.1 Runge-Kutta

In previous Monte Carlo simulations of the Zeeman decelerator [2] and magnetic trap / magneto-optical trap (MT-MOT) [151] we utilised a fourth-order Runge-Kutta (RK4) integrator [212] to solve the particle motion. We have subsequently moved to a velocity Verlet algorithm for the decelerator (see section 3.8.2) where the need to reduce computation time is very significant, since it provides adequate accuracy with fewer steps. We retain the Runge-Kutta for the MT-MOT due to its smaller global error, which is  $\mathcal{O}(\delta_t^4)$ ; the trap fields are static and so the computational overhead is smaller than in the case of the decelerator.

Consider some time-dependent function  $q(t)$  with first derivative  $\dot{q}$  given by:

$$\dot{q} = f(t, q). \quad (3.8.1)$$

If at time  $t_n$  the function has a value  $q_n$ , the RK4 method says that at time  $t_{n+1}$  the function has the approximated value:

$$q_{n+1} = q_n + \frac{\delta_t}{6}(k_1 + 2k_2 + 2k_3 + k_4), \quad (3.8.2)$$

where  $\delta_t = t_{n+1} - t_n$  is the timestep and the the variables  $k$  are:

$$k_1 = f(t_n, q_n), \quad (3.8.3)$$

$$k_2 = f\left(t_n + \frac{\delta_t}{2}, q_n + \frac{\delta_t}{2}k_1\right), \quad (3.8.4)$$

$$k_3 = f\left(t_n + \frac{\delta_t}{2}, q_n + \frac{\delta_t}{2}k_2\right), \quad (3.8.5)$$

$$k_4 = f(t_n + \delta_t, q_n + \delta_t k_3). \quad (3.8.6)$$

This method was implemented in Python using NumPy, and optimised to work with models containing  $N \geq 10^6$  particles. The atom coordinates are stored in  $N \times 6$  arrays with each row having the positions  $x, y, z$  and velocities  $v_x, v_y, v_z$ , with code methods for each of equations 3.8.1 to 3.8.6 written to work on the entire set via matrix operations, rather than using loops to iterate through them. For decelerator simulations, at each time step the code has to calculate the decelerator fields and load them into the interpolator, which is then queried to find the forces acting on the particles. The fields must then be calculated in full for the half-step time, and queried via the interpolator at

the half-step positions, and finally this process is repeated for the iterated time. Retaining the interpolator instance created at time  $t_{n+1}$  to re-use as the basis for the next loop, each iteration still requires two sets of field calculations and four interpolations. It is also necessary to make several copies of the data during operation to allow the interim calculations to be performed, which impacts code performance. Even with optimisation where possible, this is a reasonably costly method, and as will be shown in the next section faster options are available.

### 3.8.2 Velocity Verlet

In an effort to reduce computational overhead we looked for an alternative to the fourth-order Runge-Kutta algorithm. The velocity Verlet method [186] is commonly used in molecular dynamics [213], and proved to have adequate accuracy with fewer operations. In a stepwise simulation with timestep  $\delta_t$  of a particle with position  $\mathbf{r}(t)$  the following and previous positions can be written as the Taylor expansions:

$$\mathbf{r}(t + \delta_t) = \mathbf{r}(t) + \delta_t \frac{d\mathbf{r}(t)}{dt} + \frac{\delta_t^2}{2} \frac{d^2\mathbf{r}(t)}{dt^2} + \frac{\delta_t^3}{6} \frac{d^3\mathbf{r}(t)}{dt^3} + \mathcal{O}(\delta_t^4), \quad (3.8.7)$$

$$\mathbf{r}(t - \delta_t) = \mathbf{r}(t) - \delta_t \frac{d\mathbf{r}(t)}{dt} + \frac{\delta_t^2}{2} \frac{d^2\mathbf{r}(t)}{dt^2} - \frac{\delta_t^3}{6} \frac{d^3\mathbf{r}(t)}{dt^3} + \mathcal{O}(\delta_t^4). \quad (3.8.8)$$

Combining these expressions gives the standard Verlet algorithm:

$$\begin{aligned} \mathbf{r}(t + \delta_t) &= 2\mathbf{r}(t) - \mathbf{r}(t - \delta_t) + \delta_t^2 \frac{d^2\mathbf{r}(t)}{dt^2} + \mathcal{O}(\delta_t^4) \\ &= 2\mathbf{r}(t) - \mathbf{r}(t - \delta_t) + \delta_t^2 \mathbf{a}(t) + \mathcal{O}(\delta_t^4). \end{aligned} \quad (3.8.9)$$

Equation 3.8.9 can be used to solve particle motion assuming the acceleration function  $\mathbf{a}(t)$  is known for each particle, and the error term of order  $\mathcal{O}(\delta_t^4)$  can generally be neglected for small enough timesteps. However, this form is not very convenient, since it requires storing two sets of position data  $\mathbf{r}(t)$  and  $\mathbf{r}(t + \delta_t)$ . It also does not explicitly give the particle velocity, because the first-order term vanishes. The alternative ‘velocity Verlet’ method can instead be derived [214]:

$$\mathbf{r}(t + \delta_t) = \mathbf{r}(t) + \mathbf{v}(t)\delta_t + \frac{1}{2}\mathbf{a}(t)\delta_t^2, \quad (3.8.10)$$

$$\mathbf{v}(t + \delta_t) = \mathbf{v}(t) + \frac{1}{2}(\mathbf{a}(t) + \mathbf{a}(t + \delta_t))\delta_t. \quad (3.8.11)$$

Stern-Gerlach type forces arising from the interaction of a paramagnetic particle with a magnetic field depend only on position and time (if the fields are time-varying), so the method can be implemented simply; the particle position is iterated according to equation 3.8.10, the acceleration  $\mathbf{a}(t + \delta_t)$  at the new position is calculated, and then the velocity is iterated according to equation 3.8.11.

If dealing with interactions that are velocity-dependent (for example, in a laser-cooling simulation) a slightly altered formulation may be used, first calculating a half-step velocity:

$$\mathbf{v}(t + \delta_t/2) = \mathbf{v}(t) + \mathbf{a}(t)\delta_t/2, \quad (3.8.12)$$

which is used to iterate the position:

$$\mathbf{r}(t + \delta_t) = \mathbf{r}(t) + \mathbf{v}(t + \delta_t/2) \delta_t. \quad (3.8.13)$$

Finally, the acceleration at the new position is determined and used to complete the iteration of the velocity:

$$\mathbf{v}(t + \delta_t) = \mathbf{v}(t + \delta_t/2) + \mathbf{a}(t + \delta_t) \delta_t/2. \quad (3.8.14)$$

The standard Verlet method involves keeping a copy of the previous particle array in memory, which adds to computation time (the Runge-Kutta method described in section 3.8.1 has a similar drawback). Either of the velocity Verlet methods above avoid this complication, since all operations can be applied sequentially to the one array object in memory. Balanced against this is the fact that Verlet methods have a cumulative error  $\mathcal{O}(\delta_t^2)$ , compared to  $\mathcal{O}(\delta_t^4)$  in the fourth-order Runge-Kutta. Although hard to assess precisely for the decelerator, with a timestep of  $0.5 \mu\text{s}$  either method should produce negligible error, compared to other uncertainties in the system; see section 3.2. Similar assessments for the proposed MT-MOT experiment were also previously done by the author [151].

Although not strictly necessary for modelling pure magnetic trapping we have implemented the version with the half-step velocity, since it is more flexible if other (*e.g.* frictional) forces need to be incorporated into a simulation. When simulating the Zeeman decelerator, the trapping fields and the resultant acceleration on the particles are calculated as described in sections 3.6.1 and 3.4 at time  $t_c$ , which is when the coils activate. After the steps in equations 3.8.12 and 3.8.13 the field at time  $t_c + \delta_t$  is calculated and used to complete the cycle via equation 3.8.14; a key advantage over previous simulations is that the interpolator instance created at  $t_c + \delta_t$  can be stored and used as the basis for the Verlet integration at the next timestep, eliminating excess calculation. As described in section 3.7.4 the interpolator has been designed to operate efficiently on multiple queries simultaneously, with the phase-space coordinates of the particles stored in an  $N \times 6$  array. The velocity Verlet algorithm has also been implemented to operate in this way, to take advantage of the efficient matrix operations natively included in NumPy [181].

### 3.9 Point Location

‘Point location’, or the determination of which region of a partitioned space a point is located in, is a fundamental problem in computational geometry, with applications in fields such as geofencing [215] and computer graphics [216]. One of the most well-known solutions is the ‘point-in-polygon’ technique, as illustrated in figure 3.24. The simplest case to solve is a rectangular space  $R$  with minimum / maximum  $x$ -coordinates  $a/b$  and minimum / maximum  $y$ -coordinates  $c/d$ ; a point  $P$  is internal to this space if its components  $(P_x, P_y)$  lie within these bounds:

$$(P \in R) \iff (a < P_x < b) \wedge (c < P_y < d). \quad (3.9.1)$$

For a more complicated shape  $Q$  we can initially determine the smallest rectangular boundary  $R$  that contains  $Q$  and test whether our point  $P$  lies outside it by applying the inverse of the test in 3.9.1:

$$(P \notin R) \iff (P_x < e) \vee (P_x > f) \vee (P_y < g) \vee (P_y > h). \quad (3.9.2)$$

$(P \notin R) \implies (P \notin Q)$ , but  $(P \in R) \iff (P \in Q) \vee (P \notin Q)$ . In order to determine whether  $P$  lies within  $Q$  we can employ ‘ray-casting’; draw a line in any direction from  $P$ , if it crosses  $Q$  zero or an even number of times then  $P \notin Q$ , if it crosses an odd number of times then  $P \in Q$ .

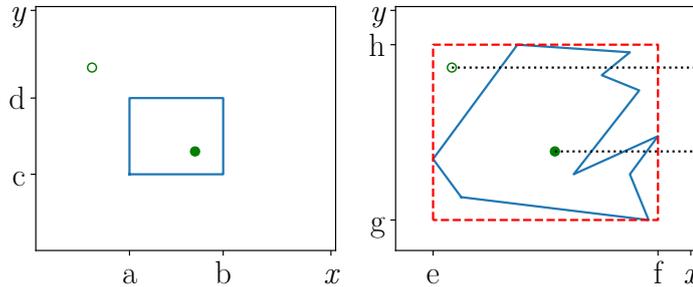


Figure 3.24: The ‘point-in-polygon’ problem in computational geometry. Left, for a rectangular region it is simple to determine whether a point lies interior or exterior to the boundary by considering whether the coordinates  $(P_x, P_y)$  lie within the bounds  $(a < P_x < b)$  and  $(c < P_y < d)$ . Right, for more complicated shapes the ray casting technique may be used. This is based on the observation that a ray drawn from the query point in any direction will cross an odd number of boundaries if it is internal, and an even (or zero) number if it is external. Example rays are shown as dotted lines. As a first approach the rectangular region technique can be applied to the smallest box that can contain the boundary (shown as a dashed line); further investigation is only needed when the query point lies within this box.

Choosing a horizontal ray pointing in the direction of positive  $x$ , we need to iteratively test whether or not it intersects each segment of the boundary. First, for segment  $IJ$  determine if the  $y$ -component of  $P$  lies between the  $y$ -components of vertices  $I$  and  $J$ :

$$[(P_y < I_y) \wedge (P_y > J_y)] \vee [(P_y > I_y) \wedge (P_y < J_y)]. \quad (3.9.3)$$

If condition 3.9.3 is not met then the ray cannot intersect the segment, and we move to the next one, otherwise we proceed as follows; if  $S$  is the point on  $IJ$  with  $S_y = P_y$  then  $S_x$  is given by:

$$S_x = I_x + (J_x - I_x) \frac{P_y - I_y}{J_y - I_y}. \quad (3.9.4)$$

If  $P_x < S_x$  then we can increment  $N$ , the number of boundary crossings. (The choice of direction of the ray is arbitrary, and we could equally well choose to test whether  $P_x > S_x$ , as long as this choice is consistent across all segments of  $Q$ ). It can be seen in figure 3.24ii that the hollow dot will have a final  $N = 2$  and the filled dot will have  $N = 3$ , indicating that they are respectively outside and inside the closed region.

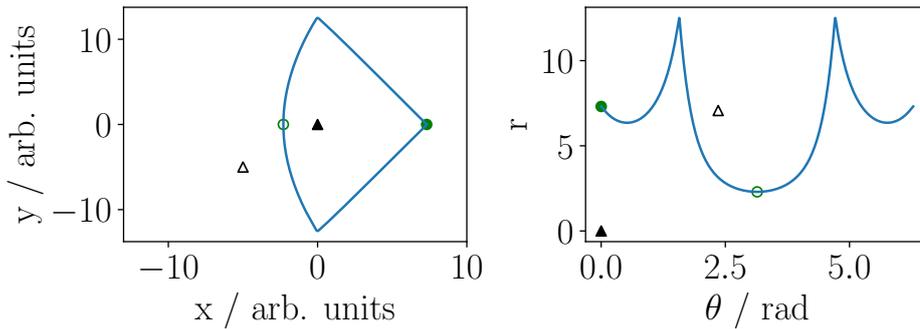


Figure 3.25: Reducing the point location problem to one easily solved via a method analogous to Monte Carlo rejection. Left, a closed surface in a Cartesian coordinate system. Filled and hollow dots mark equivalent points on the boundary, filled and hollow triangles mark points inside and outside the surface. Right, the same objects plotted in a polar coordinate system. By querying the polar boundary with a one-dimensional interpolator it is simple to determine whether a point lies above or below the boundary.

Ray-casting works well for simple polygons but is problematic for arbitrary shapes, such as the estimated separatrix (see section 4.1.2) shown on the left of figure 3.25. In order to determine whether a given point lay within this shape we could approximate it with a quadrilateral or pentagon and use ray-casting, but to improve the accuracy of the method to some arbitrary degree the boundary needs to be represented by polygons of higher order, with a corresponding computational overhead. We instead describe a simple solution combining a coordinate transformation and one-dimensional interpolation. The image on the right of 3.25 shows the boundary and points from the left image displayed in a polar coordinate system. Rather than determining whether a point  $x, y$  lies within the closed region we instead need to determine whether the corresponding point  $\theta, r$  lies below the line. The boundary is loaded into an interpolator (for example, the cubic spline (section 3.7.1) method from the SciPy library [179]), such routines are well optimised and efficient. It is then simple to test a point  $\theta, r$ ; if  $r < r_s$ , where  $r_s$  is the value the interpolator returns for  $\theta$ , then the point is within the surface. This method lends itself well to, for example, generating a set of initial phase-space coordinates within the phase-stable region of a magnetic trap, for the purposes of a simulation. In this case it is analogous to the Monte Carlo rejection technique (section 3.3), used to randomly fit data points to an arbitrary function. The method can be applied to any data known as a two-dimensional set. The use of a piecewise polynomial interpolant routine also means that the boundary does not have to be uniformly sampled, allowing shapes with a mixture of simple and complex regions to be more efficiently handled.

### 3.10 Monte Carlo Simulation of the Decelerator

This chapter has described the components that have been developed in order to enable us to simulate the travelling-wave Zeeman decelerator. The goal was to reproduce our experimental data and thus infer the characteristics of the molecular beam, and the effects of the decelerator on its velocity distribution. Once we were confident that the simulation codes were sufficiently able to capture the behaviour of the experiment, they could be used for predictive purposes, for example to allow us to rapidly determine optimal operating parameters without having to resort to trial and error, or to test new configurations or components without physically altering the apparatus. The following chapter will give examples of this predictive use, beginning with a method for calculating the power electronics delay time in section 4.1.1. This section will give an overview of how the software components that have so far been described are combined to simulate the decelerator, this will be in the specific context of guiding and deceleration of metastable argon. The experimental parameters used for the data presented in this section are given in section 2.5. All simulations were composed of  $10^6$  particles.

In order to simulate the deceleration process it is first necessary to generate a range of initial particle positions and velocities. When taking data with the decelerator we typically record a ‘transmission’ signal between every few deceleration (or guiding) events; this is simply a time-of-flight with the power electronics disengaged. As described in section 3.2 this transmission data can be fitted with a model for a supersonic expansion and the fitting parameters used in the generation of an initial sample of atoms, or we can simply fit a spline to the data and then use this to generate our source. This latter empirical method is most effective at reproducing the data although it does not allow us to give an estimate of the temperature. The settings for the power electronics are then used to calculate the time-varying currents and thus magnetic fields, these include the driving voltage, peak current, and the travelling-wave velocity. The trajectories of the particles are then solved for each timestep by interpolating the field to calculate their Zeeman interaction and numerically integrating their motion. Particles that are determined to have drifted to the sample tube are counted as lost, ones that reach a distance corresponding to the detector in the physical experiment are marked as complete and their arrival times are recorded. Once all are accounted for the simulation terminates.

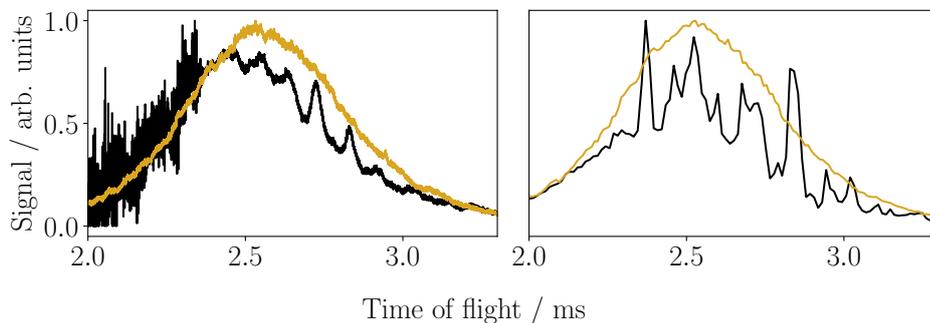


Figure 3.26: Argon deceleration, decelerator operating at  $329 \text{ m s}^{-1}$  to  $321 \text{ m s}^{-1}$ , comparison of data (left) to simulation (right). Decelerated traces in black, transmission in orange.

The arrival times can be histogrammed to give a reproduction of the experimental time-of-flight as recorded at the MCP detector. This is illustrated in figure 3.26; the left plot shows data and the right the simulation. The decelerator was set for deceleration mode from  $330 \text{ m s}^{-1}$  to  $320 \text{ m s}^{-1}$  with

a peak current of 400 A at 600 V. As described in section 4.2 the achieved wave velocity for any given set of parameters depends on how well the PWM can approximate the target waveform; for this configuration the achieved magnetic wave had initial and final velocities of 329 and 321  $\text{m s}^{-1}$ , respectively. In both cases the decelerated trace is in black and the transmission signal in orange.

Assuming the data are well reproduced by the simulation we can use the latter to infer the properties of the beam. As seen in the figure, the signal from the travelling-wave decelerator is complicated, since it is composed of overlapping signals from several adjacent and partially-loaded traps in the field. Additionally, the beam is comprised of atoms with various different values of the projection quantum number; for  $^3\text{P}_2$  argon there are five sub-levels  $m_J = -2, -1, 0, 1, 2$ , of which two are low-field-seeking, two are high-field-seeking and one is neutral. Ideally, we would expect all of the high-field-seekers to be ejected during the deceleration process, but this would still leave three possible states present in the sample. The  $m_J = 0$  atoms would be expected to produce a component to the signal equivalent to the transmission curve, since they feel no force from the field. The  $m_J = 1$  and 2 atoms would both be confined and decelerated by the field but not equally, and so would not be present in equal numbers at the detector even if they have equivalent populations at the source. This all equates to a signal with multiple peaks that is difficult to interpret. To make matters worse, some of the  $m_J = -1$  and  $-2$  atoms will certainly make it to the detector, since the decelerator is not long enough for them to all interact with the field for enough time to be ejected, and some will ‘surf’ along a local longitudinal field maximum.

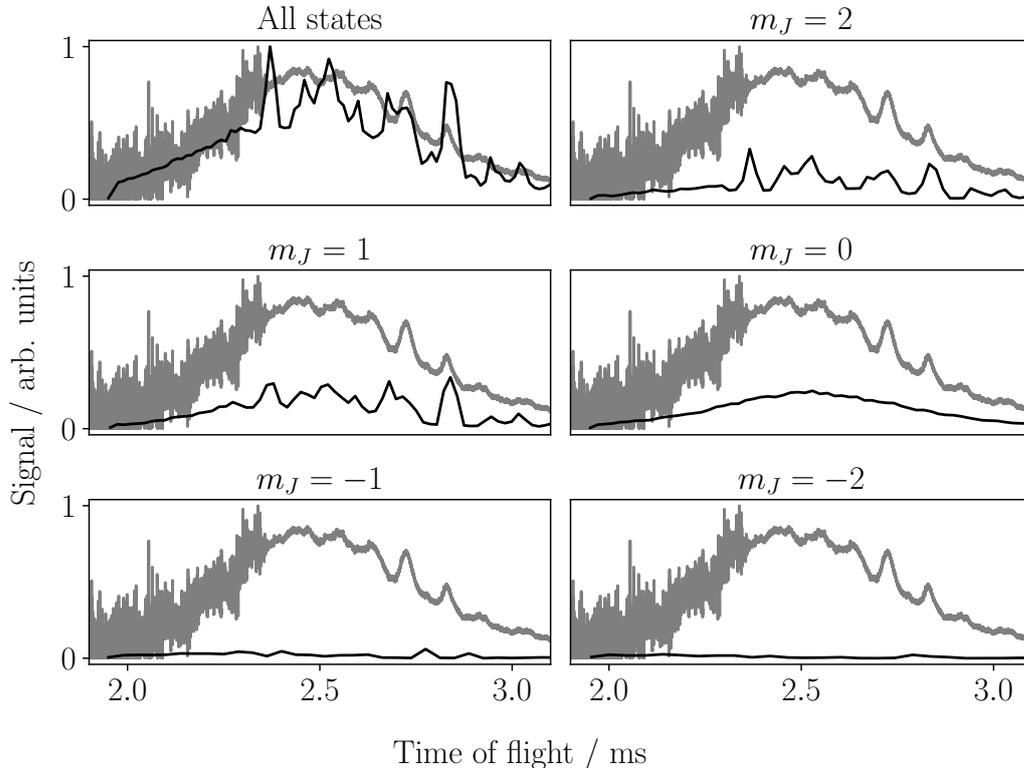


Figure 3.27: Argon deceleration, decelerator operating at 329  $\text{m s}^{-1}$  to 321  $\text{m s}^{-1}$ , comparison of data (grey) to simulation (black). Panel 1 shows the combined signal of all atoms in the simulation. Panels 2 - 6 give the individual contributions based on magnetic projection quantum number, as labelled, and are normalised to the total plot in panel 1.

We can begin to unpick the components of the time-of-flight signal by considering the contributions of the different magnetic sub-levels. Figure 3.27 shows simulations in black with the experimental trace in grey for comparison. The first panel plots the arrival times for all atoms that make it to the detector, both the real and simulated traces are scaled to their maximum value. The following panels compare the flight times for the individual sub-states, as labelled, and are normalised to the plot for all states in the first panel. It can be seen that the signal contains multiple peaks, and the simulation does a reasonable job of reproducing the main features. However, the overall shape of the total signal differs from the data, and the simulated peaks are more prominent than the real peaks.

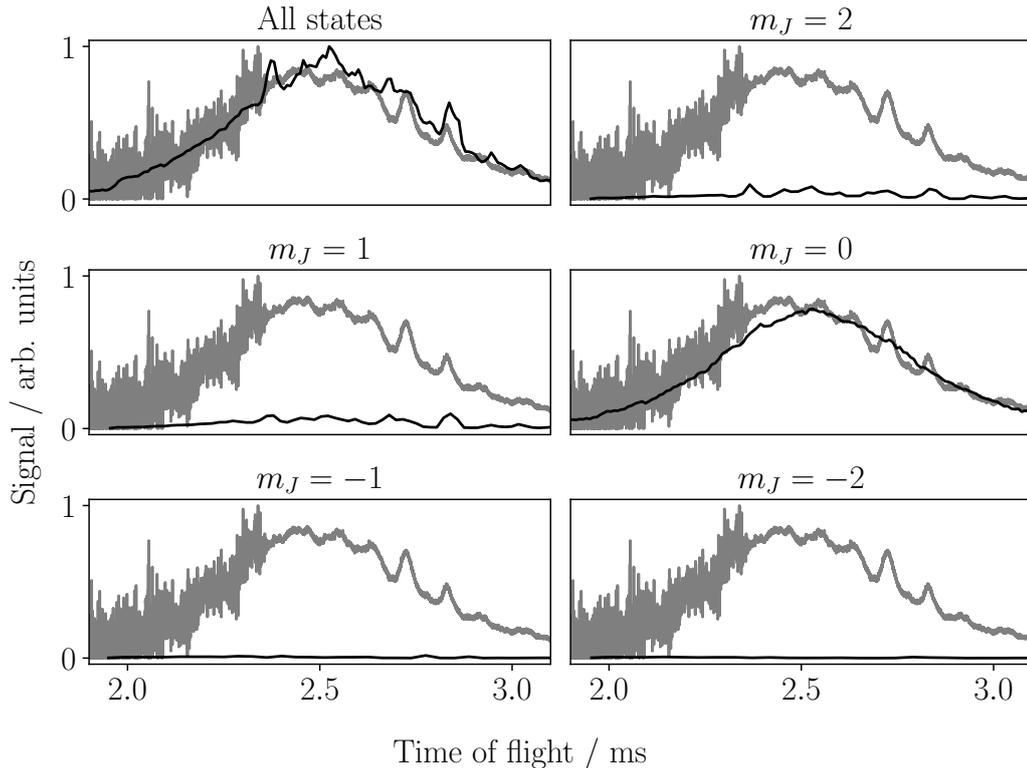


Figure 3.28: Deceleration data versus simulation, the same data as in figure 3.27 but the simulated source population has equal numbers of atoms in the  $^3P_0$  and  $^3P_2$  terms, with the latter being equally distributed among its five sublevels. This gives a much more plausible match to the data, suggesting that we are producing significant numbers of unwanted atoms in the other metastable state in our excitation source.

Whether for deceleration or guiding this pattern of stronger peaks and a differently-shaped overall signal body when compared to the data is common to all of our argon simulations. This suggests that there is a problem with our assumptions regarding the composition of states labelled by the projection number  $m_J$  in the source. As described in section 3.5, we make the reasonable assumption that for the excited sample of atoms in the  $^3P_2$  term there are roughly equal numbers of atoms for each value of  $m_J$ , which are degenerate at zero field. The population shown in the simulation in figure 3.27 is composed of approximately 33 %, 34 %, 25 %, 6 % and 3 % of  $m_J = 2 \dots -2$ , which makes sense; of the atoms surviving the deceleration process the low-field-seeking states are the most populous, followed by the neutral state, with the high-field-seekers only present in small numbers. Earlier work by the group assumed that only the  $^3P_2$  term would be present in the beam. In work on supersonic expansion of neon, Osterwalder *et al.* suggest that all three triplet states are

populated equally via electron impact excitation [217]. We are using a discharge excitation system (see section 2.2.3), however, the three argon terms lie within an energy band of only 0.18 eV, and so we may expect to see a signal composed of some mixture of  $^3P_2$  and  $^3P_0$ . We are also utilising a thermal-emission electron source (see section 2.2.4) to assist with the discharge, but we still see the disagreement between signal and simulation with the filament disengaged.

If we assume that there are equal numbers of atoms from both of the metastable terms in the supersonic expansion then instead of having 20 % of each of the five values of  $m_J$  we would have 60 %  $m_J = 0$  and 10 % of each of the others (since the term  $^3P_0$  only contains the  $m_J = 0$  sub-level). Figure 3.28 shows the result of this approach for the same data as in figure 3.27, and it can be clearly seen that this gives a much closer agreement between simulation and experiment. Practically, for a given set of experimental parameters we have found that the ratio between the populations of the two terms has to be adjusted somewhat to give the best match, but generally an even split between them works quite well.

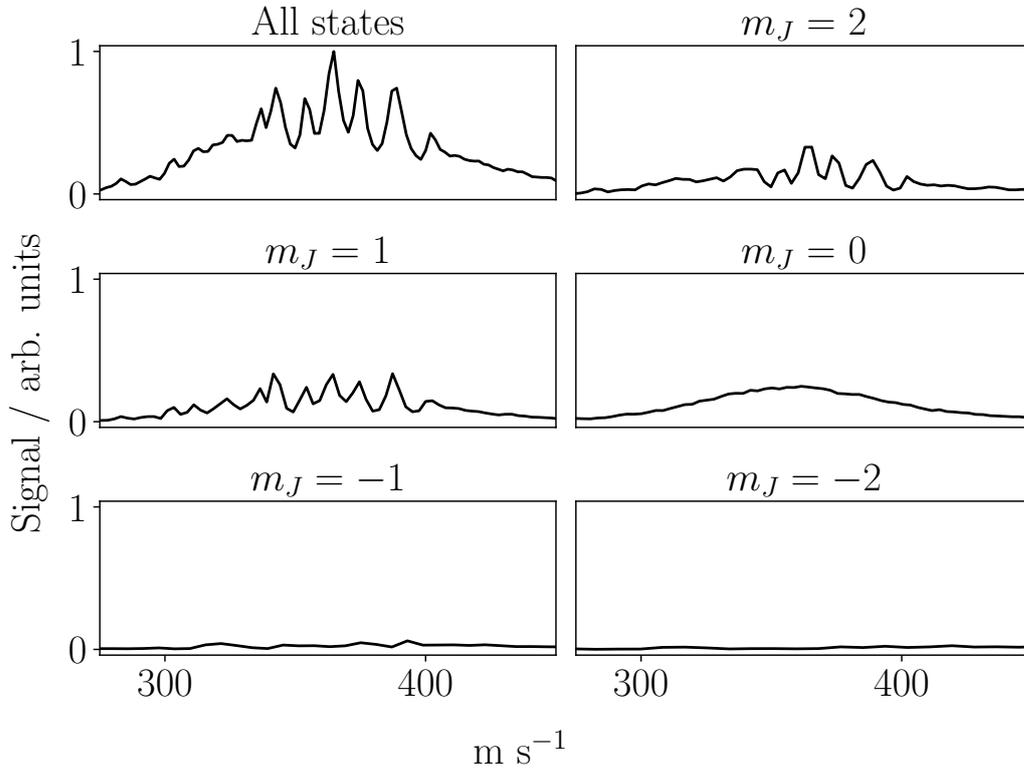


Figure 3.29: Velocities of simulated argon atoms post-deceleration, for the whole ensemble and by individual quantum state. Panels 2-6 are normalised to panel 1, giving the relative population of each state.

Being able to reproduce the time-of-flight is important, but we wish to know what effect the deceleration process has had on the velocities of the atoms. Figure 3.29 shows a comparison of the longitudinal velocities of the simulated atoms that reach the detector, in this simulation no extra  $^3P_0$  were added so each of the five sub-levels was equally represented in the initial population. As was the case for the time-of-flight plots in figure 3.27 we see significant manipulated peaks in the low-

field-seeking states and depletion of the high-field-seekers, with the  $m_J = 0$  atoms unaffected by the decelerator fields. A closer view is given in figure 3.30, which compares the complete simulated ensemble and the low-field-seekers, all normalised to the same scale. Also shown is a velocity distribution generated from experimental data, using the method outlined in section 3.2, but this must be treated with caution. Recovering a velocity distribution from a time-of-flight profile is justifiable if we can assume a straightforward correlation between flight time and velocity; this is the case for a free-flying molecular beam which is not undergoing collisions or subject to any external forces, as long as the pulse is very short, *i.e.* we neglect the temporal (and spatial) widths of the source. Although the bulk of the plot represents atoms in a non-paramagnetic state, the peaks in the data do not. Both high- and low-field-seeking atoms may have been accelerated or decelerated variously during their journey through the decelerator and we cannot assume that the effects average out. Regardless, there is reasonable agreement between the top two panes in the figure. The red and blue dashed lines indicate the initial and final velocities achieved by the moving trap. There are small peaks for  $m_J = 1$  and 2, although surprisingly the latter is weaker than the former. This will be explored shortly.

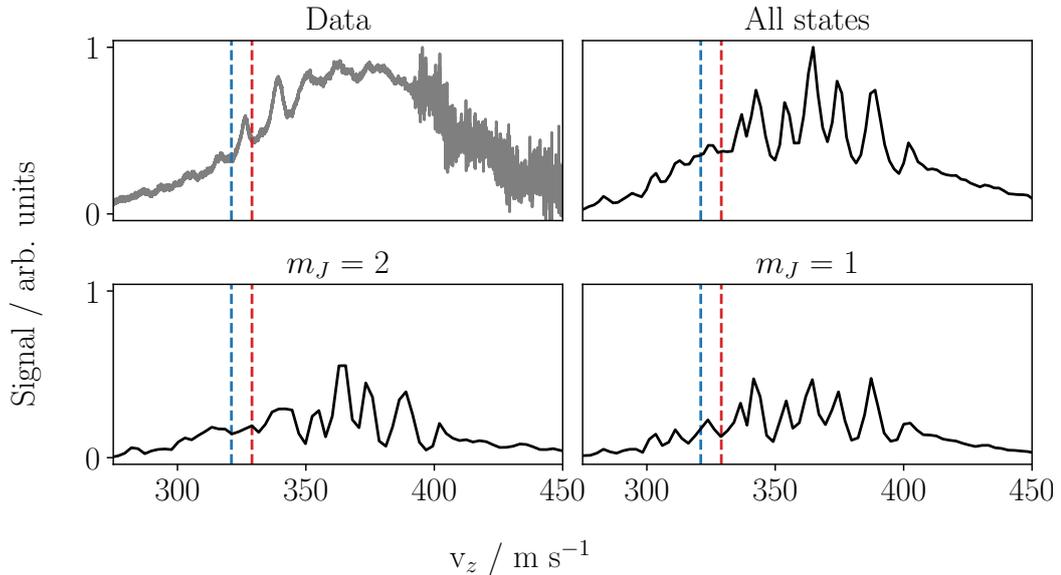


Figure 3.30: Time-of-flight data converted to a velocity profile, compared to the velocities in the simulation for all magnetic states and the two low-field-seeking states. Red / blue lines indicate the initial and final velocities of the decelerator wave.

We know the velocity distribution that we have assigned to the complete ensemble of particles at the source, but comparing this distribution to the velocities of the atoms that survive to reach the detector does not tell us very much. More useful would be to know what the initial conditions of these specific atoms were, rather than the whole set. In order to achieve this the simulation code saves a copy of the source distribution at the instant at which it is fully populated. Each atom has a unique index so that the final sample at the detector can be related back to the beginning. Figure 3.31 shows the initial (red) versus final (blue) longitudinal velocities for all atoms in the sample, and the same for only the  $m_J = 2$  atoms. The peak visible between the lines in figure 3.30 can be seen, and appears to correspond to a peak in the source that has been shifted. In order to better determine what has happened we can consider also the physical positions of the atoms in the decelerator wave.

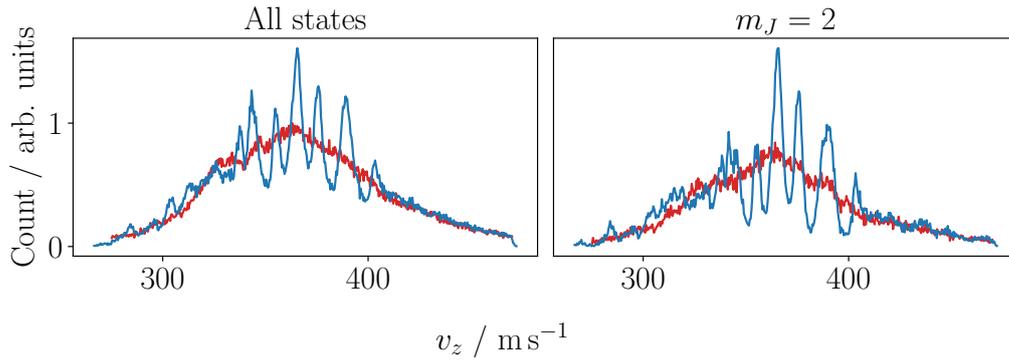


Figure 3.31: Comparison of initial (red) to final (blue) velocity distributions for only the argon atoms that survive to reach the detector. Left, all 5 sub-levels, right, only the spin-stretched  $m_J = 2$  atoms. There appears to be a shift of atoms to the left in the region of the velocity of the decelerator trap, indicating a deceleration.

The decelerator field sequence can be calculated and inspected to determine the position as a function of time of the target individual trap in the travelling wave that we aim to load, as discussed previously. This gives us the times at which the trap is fully formed and about to collapse, and the simulation code takes a snapshot of the particle positions and velocities at these times. These can be plotted against the decelerator fields as a longitudinal phase-space plot, shown in figure 3.32; left is the atom sample and field at the time the trap forms, right at the time it is about to shut down. For the experimental parameters we are currently considering these times  $t$  are at  $t = 0.9 \text{ ms}$  and  $t = 2.23 \text{ ms}$  after the beam source is activated. Here we are considering the second trap to form in the wave, marked by the dashed green line. This is the trap we typically aim to load; historically this was due to the alternating focussing / defocussing interaction of the wire quadrupole with the decelerator field (see section 2.3.7) although the data shown here had no external quadrupole applied. The trailing trap is marked in orange.

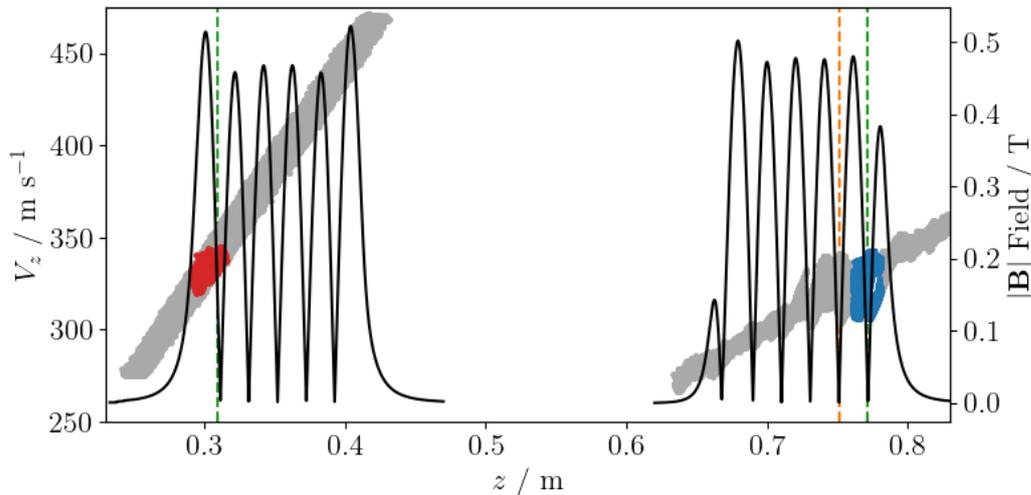


Figure 3.32: Longitudinal phase-space plots of argon before and after deceleration. The positions and velocities of the atoms are saved at the instant at which the second trap in the wave is fully formed (0.9 ms, left plot), and just as it is about to collapse (2.23 ms, right plot). The trap centre is marked by the green line; trap 3 is marked in orange, see figure 3.35. The atoms located within the trap at the end are highlighted in blue, and their positions at the start in red.

The atoms that are between the maxima of the trap at the end of the sequence are marked in blue, and we can identify their positions and velocities in the earlier sample, shown in red. Figure 3.33 shows only the  $m_J = 2$  atoms. Comparing the left and right plots we can see that not all of the atoms that were present in trap 2 at the beginning are there at the end; some will have been lost to collisions with the tube wall, others will have escaped to a neighbouring trap. The large thermal spread of longitudinal velocities in the beam is seen, and is much larger than the acceptance of the decelerator, so that the majority of the atoms will not be phase-stable at any point in the magnetic wave and will eventually be ejected. However, for the four-module decelerator there is not enough time for this process to complete, and it can be seen in the blue plot that some atoms from the trailing trap are in the process of entering the target trap. The phase-stable region of the trap is not uniformly filled but given enough time the trapped atoms would rotate in phase-space and a somewhat more even distribution would form, with the atoms orbiting around the centre of the potential.

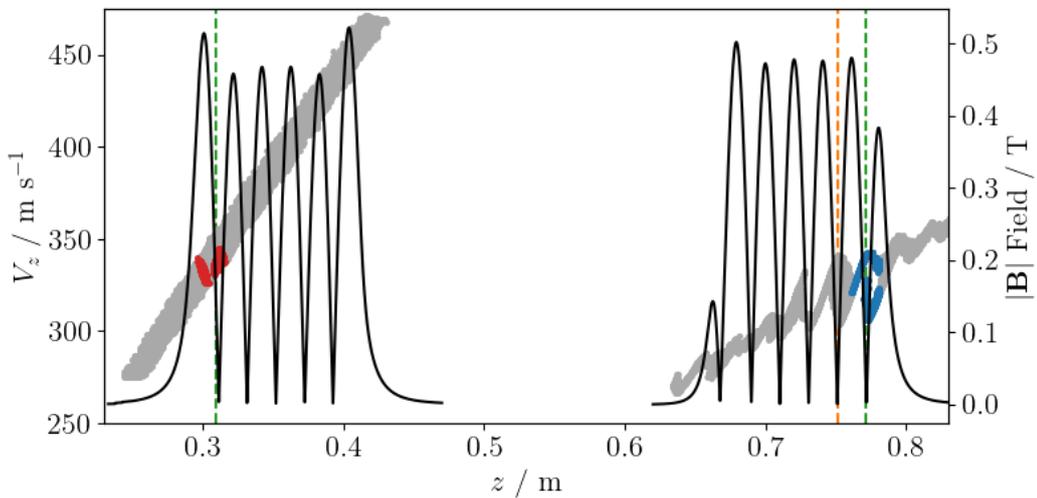


Figure 3.33: As per figure 3.32, but only the atoms in the  $m_J = 2$  state are shown.

The shift in velocities for only the atoms located in trap 2 at the end of the deceleration sequence is shown in figure 3.34. We are most interested in the spin-stretched  $m_J = 2$  atoms shown in the second panel, which appear to show significant deceleration; the peaks in this plot can be identified in figure 3.31. Analysis of these atoms shows that their mean velocity has decreased from  $336 \text{ m s}^{-1}$  to  $325 \text{ m s}^{-1}$  during their passage through the decelerator.

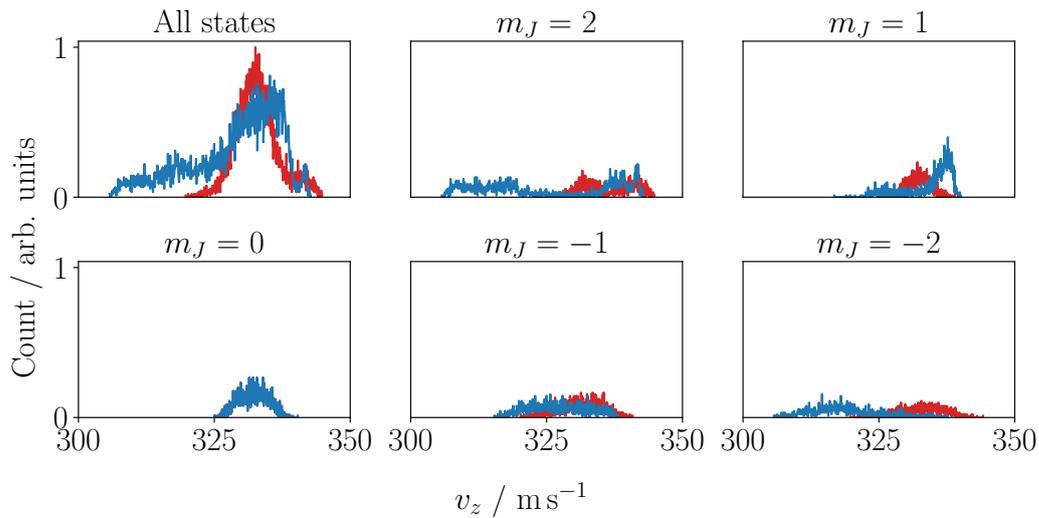


Figure 3.34: Velocities of the simulated argon atoms that remain in trap 2 at the end of the deceleration sequence, as illustrated in figure 3.32. Red / blue are initial / final velocities. There is not enough time for all phase-unstable particles to be ejected but a deceleration of the spin-stretched  $m_J = 2$  atoms is seen.

Also visible in figure 3.33 is a bunched set of atoms in trap 3 (marked by the orange line), which is trailing the trap containing the set marked in blue. For the current set of decelerator parameters this trap forms at  $t = 0.96 \mu\text{s}$  after the source triggers and begins to collapse at  $t = 2.34 \mu\text{s}$ . Following the same procedure as described for trap 2 we can save the atom distributions at these times and examine them to determine what effect the deceleration process has had on the atoms that end the sequence in trap 3. This is shown in plot 3.35. The shapes of the initial and final phase-space distributions of the trap 3 atoms can be seen to be somewhat different to those of the trap 2 atoms; at the beginning of the sequence the trap 3 atoms are already loaded into the trap and stay in it until the end. They also form a distinct shape that suggests that the longitudinal separatrix of this trap is reasonably well-filled compared to the previous example. However, when we look at the velocity distributions for the components, shown in figure 3.36, we do not see a clear example of deceleration as we did for trap 2. It seems that the part of the beam loaded into this trap was centred around a velocity slightly lower than that initially possessed by the wave. The mean longitudinal velocity of the  $m_J = 2$  atoms is practically unchanged, going from  $324 \text{ m s}^{-1}$  to  $323 \text{ m s}^{-1}$ . This illustrates the care that must be taken when interpreting the results of this experiment, especially when such a modest deceleration as the one shown here is used.

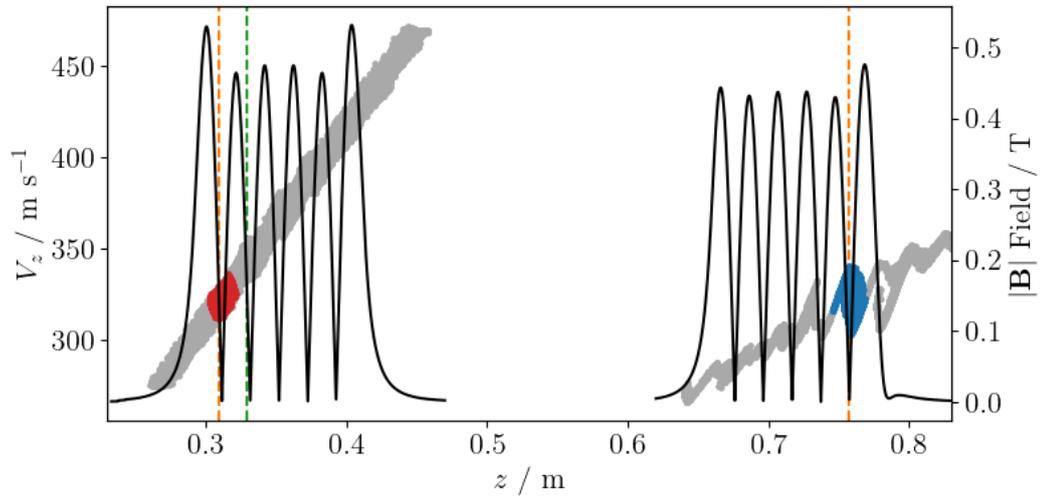


Figure 3.35: As per figure 3.33 but considering the atoms in trap 3, marked by the orange line, compared to trap 2, green line. Right, the trap is about to collapse at 2.34 ms after the source triggered. Atoms within the trap are in blue. Left, the trap has just formed at 0.96 ms. The atoms that will be in the trap at the end are in red. The positions and velocities of the atoms entering this trap are not well matched to it and so little deceleration is seen.

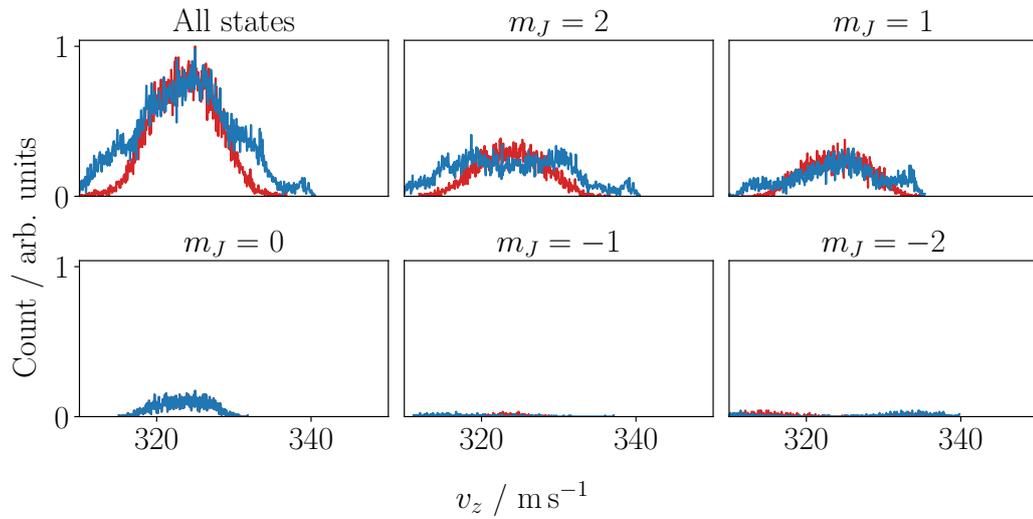


Figure 3.36: Velocities of the simulated argon atoms that remain in trap 3 at the end of the deceleration sequence, as illustrated in figure 3.35. Red / blue are initial / final velocities.

Figure 3.37 shows pre- and post-deceleration of the atoms in trap 2, but this time for the  $m_J = 1$  state. Compared to figure 3.33 it can be seen that the range of longitudinal velocities possessed by the atoms is significantly smaller, which is to be expected; they are less strongly trapped, and so their phase-stable region is reduced. This appears to be relevant to the somewhat surprising results shown in figure 3.30, in which  $m_J = 1$  has a sharper, narrower peak in the region of the target velocity than  $m_J = 2$ .

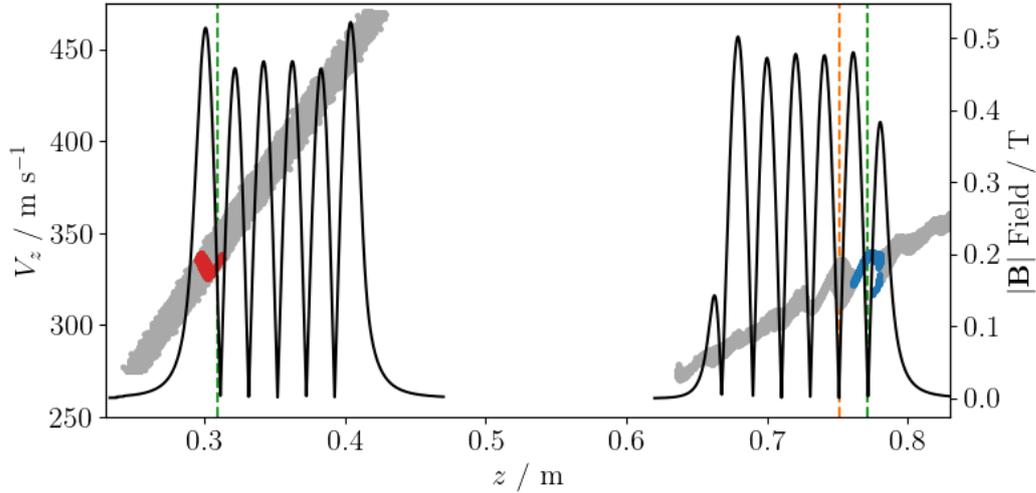


Figure 3.37: As per figure 3.33, but showing the atoms in the  $m_J = 1$  state.

Remaining on the subject of the distribution of low-field-seeking states in this simulation; it was mentioned that both  $m_J = 1$  and  $m_J = 2$  are present in roughly equal numbers in this simulation, presumably because there is not enough time for the less strongly trapped particles to be lost. We should therefore predict that a longer decelerator would produce a final sample population with a greater prevalence of  $m_J = 2$  than  $m_J = 1$ , although the latter should have a smaller range of velocities. We would also expect fewer  $m_J = 0$ , since they would have more time to be lost to lateral drift, and few to no high-field-seeking atoms. These predictions are supported by the results of a simulation with the same initial conditions as the work shown so far in this section, but with 9 instead of 4 decelerator modules; at the end of the deceleration sequence in the longer machine there were 4 times as many  $m_J = 2$  as  $m_J = 1$ . Further details are given in section 4.5.1.

For comparison, an example of constant-velocity guiding mode is shown in figures 3.38, 3.39 and 3.40. The decelerator was configured to operate at  $330 \text{ m s}^{-1}$  at 400 A; the PWM achieved this with negligible error. Other parameters were the same as for the previous deceleration case. A clear enhancement of the signal around the target velocity can be seen. The  $m_J = 2$  atoms in the second trap, shown in figure 3.40, had mean initial / final velocities of  $339 / 327 \text{ m s}^{-1}$ .

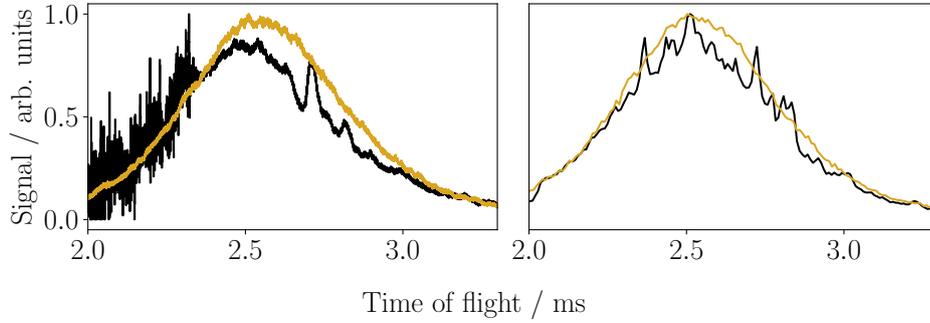


Figure 3.38: Argon guiding, decelerator operating at  $330 \text{ m s}^{-1}$ , comparison of data (left) to simulation (right). Guided traces in black, transmission in orange. As described in the main text, extra atoms in the  $m_J = 0$  state have been added to the simulation to account for the suspected presence of  $^3\text{P}_0$  in the excitation source.

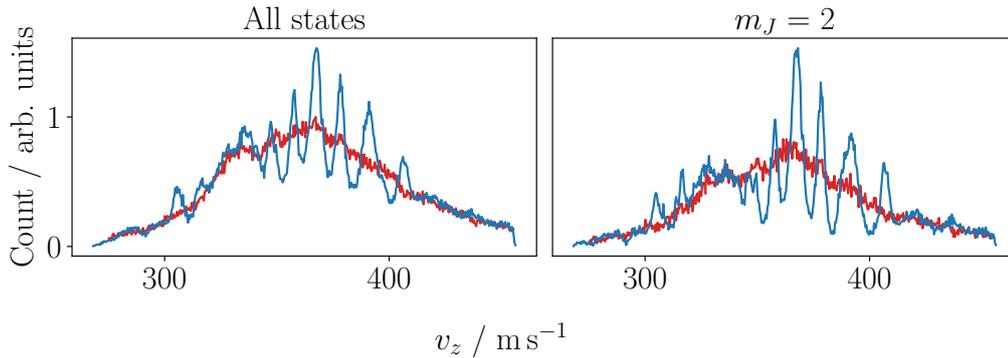


Figure 3.39: Guiding mode, comparison of initial (red) to final (blue) velocity distributions for only the argon atoms that survive to reach the detector. Left, all 5 sub-levels, right, only the spin-stretched  $m_J = 2$  atoms.

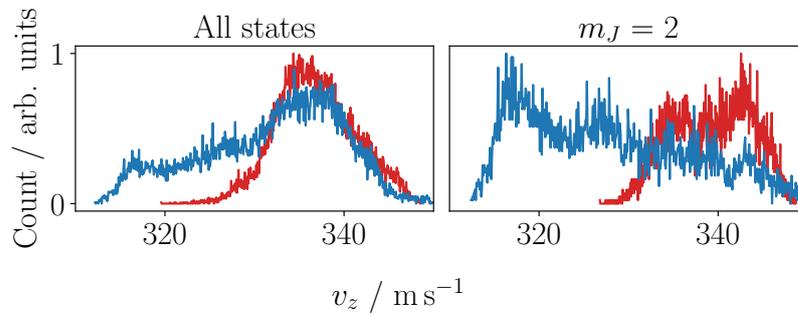


Figure 3.40: Velocities of the simulated argon atoms that remain in trap 2 at the end of the guiding sequence. Red / blue are initial / final velocities. At the end of the sequence the  $m_J = 2$  atoms in this trap had a mean velocity of  $327 \text{ m s}^{-1}$ , compared to the trap velocity of  $330 \text{ m s}^{-1}$ .

The examples shown so far have compared simulations with data for argon guiding and deceleration, with a target initial trap velocity of  $330 \text{ m s}^{-1}$ . Establishing the beam flow velocity by analysing the transmission time-of-flight signal was discussed in sections 3.1.1 and 3.2, which gave alternative models in equations 3.1.8 and 3.2.7. Fitting these models to the data in this section shown so far gives the flow velocity of the beam as being approximately  $350$  to  $360 \text{ m s}^{-1}$ ; this seems plausible when compared to the velocity plots in figures 3.30, 3.31 and 3.39. The trap is therefore moving more slowly than the bulk of the atoms in the beam. An alternative approach is to try and match the travelling-wave to the central flow velocity of the beam, maximising the number of atoms available to be captured.

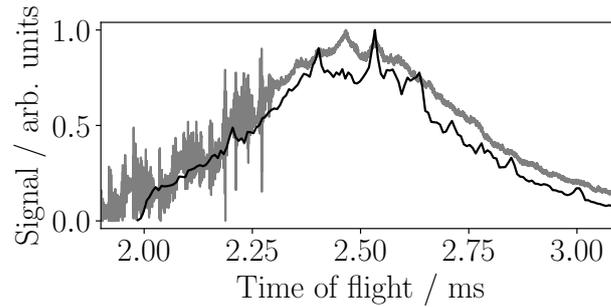


Figure 3.41: Deceleration of argon, trap initial / final velocity of  $361 / 350 \text{ m s}^{-1}$ , peak current of  $400 \text{ A}$ . Experimental data is in grey with simulated time-of-flight overlaid in black, the latter padded with additional  $^3\text{P}_0$ . When attempting to capture the central portion of the beam the manipulated peaks are obscured by the bulk signal.

Figure 3.41 shows such an attempt. The beam has the same flow velocity and temperature as the previous examples but the decelerator was configured with initial and final velocities of  $360$  and  $350 \text{ m s}^{-1}$  at  $400 \text{ A}$ ; analysis of the simulation shows that the PWM synthesis actually achieved  $361$  to  $350 \text{ m s}^{-1}$ . Additionally, the experimental runs with an initial trap velocity of  $330 \text{ m s}^{-1}$  were configured with a power electronics delay of  $796 \mu\text{s}$ ; this turned out to be a sub-optimal setting, as will be described in section 4.1.1. The data in figure 3.41 had a delay of  $749 \mu\text{s}$ , which was better chosen to match the central velocity of the atom sample into the centre of the phase-stable region of the moving trap.

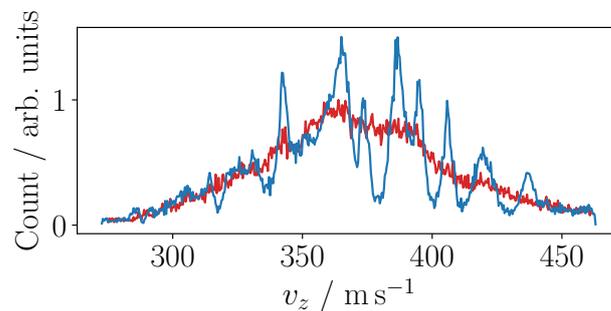


Figure 3.42: Simulation shown in figure 3.41, comparison of initial (red) to final (blue) velocity distributions for only the  $m_J = 2$  atoms that survive to reach the detector.

Extra  $^3\text{P}_0$  atoms were added to the simulation, in this case a ratio of 2:1 of  $^3\text{P}_0$  to  $^3\text{P}_2$  was necessary to reproduce the bulk of the data. Any peaks in the fast wing of the data are as usual obscured by interference from the decelerator power electronics, those in the slower wing appear to be hidden by the high background of non-paramagnetic  $^3\text{P}_0$  atoms. The central peaks are well represented, however, and they are the ones that are relevant.

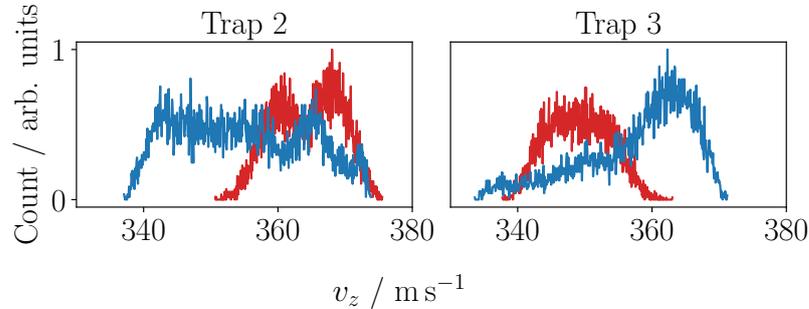


Figure 3.43: Initial (red) and final (blue) velocities of the simulated  $m_J = 2$  argon atoms that remain in trap 2 (left) and trap 3 (right) at the end of the deceleration sequence. In trap 2 the mean velocity decreased from 364 to 355  $\text{m s}^{-1}$ , in trap 3 it increased from 349 to 358  $\text{m s}^{-1}$ .

Considering only the atoms that were located within the maxima of the second trap in the wave (which was the loading target), we see a clear shift in the longitudinal velocities, as shown in the left plot of figure 3.43. The initial / final velocities are as usual depicted in red / blue, the mean value of which decreased from 364  $\text{m s}^{-1}$  at the instant the trap formed to 355  $\text{m s}^{-1}$  at the time it was about to collapse. The right plot shows trap 3, and it can be seen that the fraction of the beam it captured was on average slower than the initial trap velocity, so that some of the atoms were actually accelerated. The end result is that both traps produced spatially bunched samples of atoms with velocities in the region of the target velocity. By the time they reached the detector these bunches would expand and partially merge, giving rise to the multi-peak signal seen in figure 3.41. Admittedly, these peaks are not very prominent, given the number of non-paramagnetic atoms obscuring them.

We will conclude this section with a recap of the chapter. We have demonstrated the application of the set of computer codes that we have developed in order to simulate the travelling-wave Zeeman decelerator, with details of its individual components. We have attempted to justify the choices that were made regarding the physical models that were used, and the computational methods used to achieve our goals. These simulation codes represent the main result of the work described in this thesis, and have granted us significantly greater insight into the operation of the physical experiment than we previously had. In addition to their application to our decelerator experiment, the individual code libraries that we have made available for use by others have contributed to multiple research projects, not only within our group but also at other institutions.

This section has compared some experimental data with reproductions of those data, for constant-velocity guiding and deceleration of metastable argon. Firstly, we have shown that we are able to match the major structure of the time-of-flight signal, as seen in figures 3.26, 3.38 and 3.41. Although there is some variation in their relative intensities, the temporal positions of the manipulated peaks is reproduced in the simulations. The quality of the time-of-flight signals produced by the simulations gives us confidence that the velocity distributions they provide can give some good

quantitative information on the actual velocities of the atoms in the molecular beam. The problem of matching flight times to velocities in a supersonic expansion was described in chapter 2 and in sections 3.1.1 and 3.2; we have argued that for our experiment we can neglect complications arising from the spatial and temporal widths of the source and we can reliably reproduce the distribution of molecular velocities in our experimental data by fitting either a ‘real’ model based on physical principles or simply an empirical model to match the data. This is much more complicated when the interaction of the particles with the decelerator fields is included and there is no longer a relatively simple correlation between arrival time and velocity. Our ability to produce fully time-dependent simulations of the dynamical interactions of the beam with the decelerator has been key to addressing this problem, which was only partially addressed in earlier work that had to rely on semi-dynamical models.

The signals produced by the experiment when guiding or decelerating are much more complicated than the relatively simple skewed Gaussian obtained from zero-field transmission. This is due to several factors; the decelerator field is comprised of multiple sequential traps, which are differentially loaded with parts of the supersonic expansion. Even though the Even-Lavie valve produces very short pulses, the thermal expansion of the beam means that it rapidly spreads out along the propagation axis. The atoms in each trap may or may not be phase-stable, and will either escape to the leading or trailing traps. Once the decelerator field collapses even groups of atoms from individual traps that are well bunched spatially will expand and can partially merge with each other, making the interpretation of the signal received at the detector more difficult. The presence of multiple magnetic sub-levels in the case of most species of interest confounds matters further, since they will not all be equally manipulated by the field.

All of this needs to be accounted for in the codes in order to fully reproduce the signal; once that is achieved we can isolate the parts of the simulation result that we are interested in, confident that it is representative of the data. This is shown in *e.g.* figures 3.31 and 3.34, where we examine the pre- and post-deceleration velocities of the entire beam, or only those contained within a single trap. Earlier models of the experiment generated a time-average of this one trap and were able to make some analysis this way, but we have shown that in order to match the experimental data the entire beam needs to be included. This has allowed us to attempt to quantify the operation of the decelerator in a detailed way for the first time. Two approaches have been shown in this section; attempting to match the moving trap to the central velocity of the beam, where the particle number density is highest, or instead trying to capture a slower-moving part of the beam, so that the acceptance is greater for an equal reduction in velocity. The data available were not optimal, but regardless there is good evidence that the decelerator is operating as intended. Our simulation codes now allow us to test various sets of operational parameters and assess their performance. We have extended them to allow us to also test modifications to the experimental apparatus, such as increasing the length of the deceleration stage or changing the transverse-focussing quadrupole guide. We can also estimate the acceptance and maximum deceleration for other paramagnetic species, allowing us to compare our results with those of other travelling-wave Zeeman decelerator experiments. These topics will be the subject of the following chapter.

# Chapter 4

## Applying the Simulations

### 4.1 Decelerator Loading

#### 4.1.1 Delay Calculation

When the power electronics of the travelling-wave Zeeman decelerator are activated with a programmed pulse sequence, a series of longitudinal magnetic traps are created, which move in turn along the length of the decelerator. In order to successfully load part of the molecular beam into one of these traps the sequence needs to be activated an appropriate time after the pulsed valve has triggered. This delay time can be estimated by borrowing a concept from the field of particle accelerators; the ‘synchronous atom’, an idealised particle that experiences the optimum manipulation from the driving fields. The following example is for the specific case of the low-field-seeking  $m_J = 2$  state of  $^3\text{P}_2$  metastable argon.

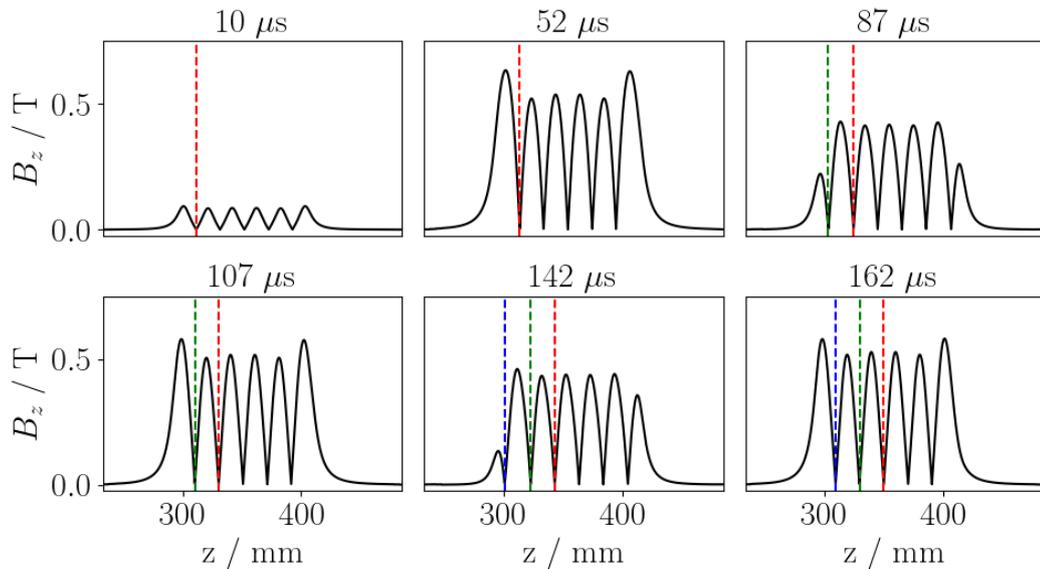


Figure 4.1: Sequence showing the longitudinal decelerator field, operating in guiding mode at  $336 \text{ m s}^{-1}$ , 600 V and 400 A, the wave ‘travels’ from left to right. Top, left; 10  $\mu\text{s}$  after coil activation, trap 1 is forming (red dashed line) - this is the first trap which could carry an atom through the entire deceleration sequence, and so the traps to the right of it are not considered. Centre; trap 1 is fully formed. Right; trap 2 is forming (green line). Bottom, left; trap 2 is fully formed. Centre; trap 3 is forming (blue line). Right; trap 3 is fully formed.

We consider the trajectory of an ideal atom, with creation time and start position midway in the range of possible values. Excited argon atoms are created in a time interval determined by the frequency and number of pulses of the dielectric barrier discharge (DBD, see section 2.2.3), typically of order 20 to 30  $\mu\text{s}$ . The central point of the DBD is at  $z = 6 \text{ mm}$ . The velocity distribution of the atoms

is estimated by fitting a model of supersonic expansion (see section 3.2) to an experimentally determined time-of-flight; for argon produced in our Even-Lavie valve cooled to 140 K the gas packet has a temperature of  $\approx 6$  K and a typical velocity in the lab frame of  $\approx 360 \text{ m s}^{-1}$ . Ideally, the decelerator will be programmed with a travelling wave matching the central velocity of the beam, but if the beam is too ‘hot’ then it might not be possible for a valid pulse sequence to be generated (see section 2.3.6). In this case a lower frequency wave may be chosen to capture a slower portion of the beam.

Figure 4.1 illustrates the first three traps forming for a  $336 \text{ m s}^{-1}$  constant velocity guiding sequence, with maximum coil current of 400 A driven by 600 V. As a first attempt, we considered an arrangement such that the ideal atom arrives at the centre of trap 2 just as the rear wall of the trap has risen to 50% of its full value. As shown in panel 3 this occurs  $87 \mu\text{s}$  after the power electronics activate; the trap centre, marked by the green dashed line, is at  $z = 303 \text{ mm}$ . Historically, we aimed to load trap 2 in the decelerator, because odd-numbered traps in the field were defocussing when combined with the wire quadrupole, see section 2.3.7. This problem would not occur with the permanent magnet quadrupole. Choosing a typical set of parameters for the DBD of 18 pulses at 550 kHz gives a discharge time of  $32.7 \mu\text{s}$ , so the ideal atom is considered to be created in the source apparatus at  $t_c = 16.4 \mu\text{s}$ . The estimated time delay for the power electronics  $t_d$  is given by  $t_d = t_a - t_r$ , where  $t_a$  is the arrival time of the atom and  $t_r = 87 \mu\text{s}$  is the time elapsed since the coils activated (note that this includes a  $3 \mu\text{s}$  overall lag in the control system for the power supplies, see section 2.3.5). Neglecting an interaction with the magnetic fields the flight time of the ideal atom is  $t_f = (303 - 6) \times 10^{-3} \text{ m} / 336 \text{ m s}^{-1} = 8.839 \times 10^{-4} \text{ s}$ . Arrival time  $t_a = t_c + t_f \approx 900 \mu\text{s}$ , giving a delay time of  $813 \mu\text{s}$ . The same technique can be used to determine the delay for an ideal atom arriving at a different point in the decelerator sequence by adjusting the values of  $t$  and  $t_r$ ; for the current example better results were obtained by aiming to load the ideal atom into trap 2 as it is fully formed at  $t = 107 \mu\text{s}$ , as in panel 4 of figure 4.1.

The interaction between the ideal atom and the rising decelerator field was added to the model using the method detailed in section 3.4. The code creates the ideal atom with characteristics as described, and begins modelling its motion through the decelerator with an arbitrary initial electronics delay  $t_d = 500 \mu\text{s}$ . Once the atom reaches the point  $z = 303 \text{ mm}$  the simulation stops. The difference between the actual and target arrival times is  $\delta t = t_a - (t_d + t_r)$ ; if  $\delta t$  is too large, say  $|\delta t| > 1 \mu\text{s}$ , then the sequence is restarted with an iterated delay  $t'_d = t_d + \delta t$ . The trajectory simulation runs with a timestep of  $0.5 \mu\text{s}$  so two loops is generally enough to obtain an adequate estimate for the delay time, as illustrated in figure 4.2. For the example given here of argon moving at  $v_z \approx 336 \text{ m s}^{-1}$  the effect of the decelerator field on the arrival time at the forming trap is found to be negligible, as seen in figure 4.3; for the range of plausible mean beam velocities achievable with a pulsed source of cryogenic argon the error in the estimated optimal loading delay is no more than  $0.4 \mu\text{s}$ . This is not necessarily the case for much slower beams, or species with lower mass / higher magnetic moment. Figure 4.4 shows a comparison of a free-flight loading model and one including the field interaction for hydrogen; at  $50 \text{ m s}^{-1}$  the free-flying ideal hydrogen atom will arrive  $25 \text{ m s}^{-1}$  too early, significantly affecting the trap-loading efficiency.

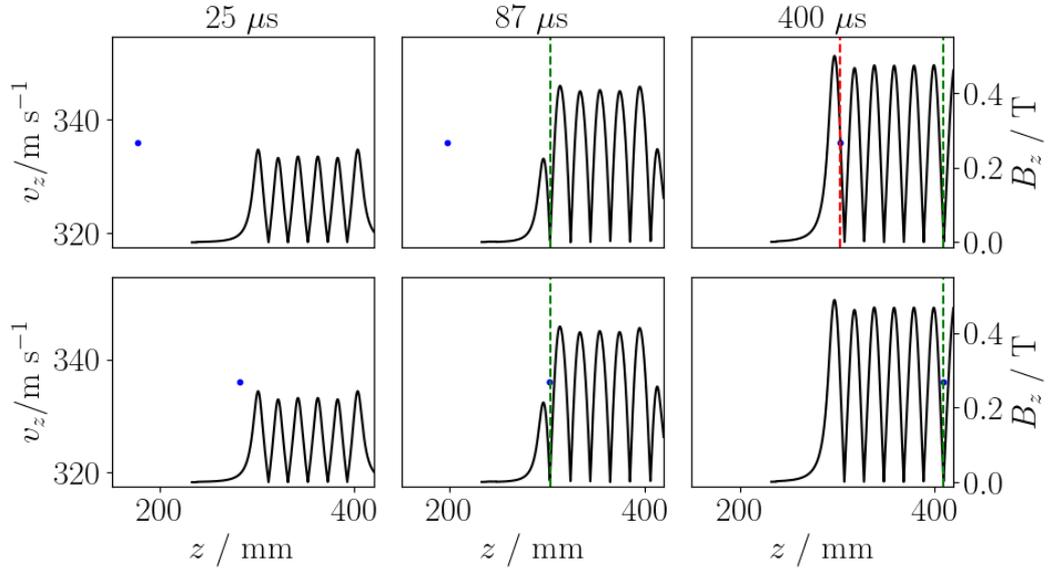


Figure 4.2: The operation of the software model used to estimate the time delay parameter for the power electronics in order to load an ‘ideal’ atom into the moving trap. Top, arbitrary delay of  $500\ \mu\text{s}$ , left;  $25\ \mu\text{s}$  after coil activation. Centre; the field has progressed to the loading point, the green dashed line indicates the trap minimum. Right; the atom arrives at the target location (red dashed line) too late, and is loaded into the wrong trap. The position it should be in is shown by the green line. The code measures the difference between target and achieved arrival time and adds this to the delay, and restarts. Bottom, delay of  $813\ \mu\text{s}$ , left;  $25\ \mu\text{s}$  after coil activation. Centre; the atom arrives at the correct trap minimum, the difference between the target and arrival times is within  $1\ \mu\text{s}$  so the software successfully terminates. Right; comparison with top right of situation  $400\ \mu\text{s}$  after switch-on, with the atom correctly loaded.

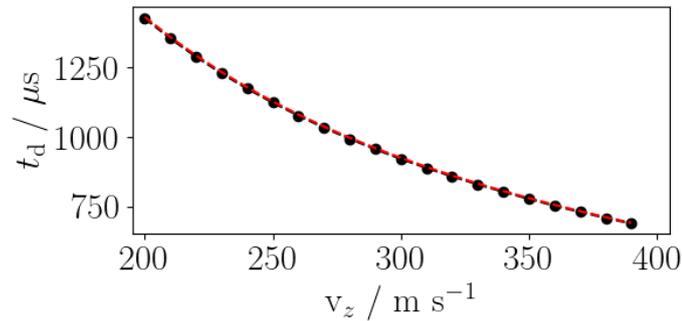


Figure 4.3: Power electronics delay time required to deliver an ideal on-axis particle into the decelerator field at the optimal time, as a function of particle velocity. The dots show results of an iterative trajectory simulation including the interaction with the decelerator field, for  $^3\text{P}_2$  argon. The dashed red line indicates the delay neglecting this interaction; the difference between the two values is no more than  $0.4\ \mu\text{s}$  across the range of interest.

The estimated delay value is a good starting point when setting up the decelerator, but the three-dimensional shape of the traps in the field is complicated and so it is not surprising that simply aiming the central atom in the packet at the centre of the trap does not necessarily give optimal results; in practice further work must be done to optimise the signal at the detector.

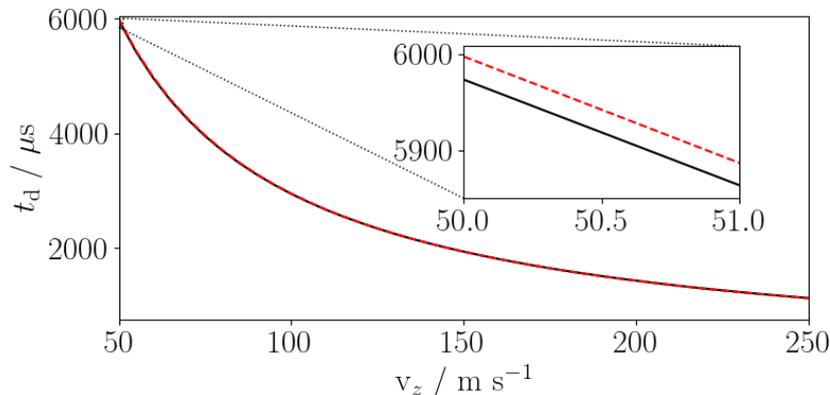


Figure 4.4: Power electronics delay time required to deliver an ideal on-axis particle into the decelerator field at the optimal time as a function of particle velocity for (low-field-seeking)  $^2S_{1/2}$  hydrogen. The black curve shows the results of an iterative trajectory simulation including the interaction with the decelerator field, the dashed red line is a simple time-of-flight model neglecting this interaction. At higher velocities the difference between the two values is negligible, as for the argon case, however for slower beams the estimated delay is wrong by  $\approx 25 \mu\text{s}$  (shown in inset).

### 4.1.2 Longitudinal Acceptance Model

Section 4.1.1 discussed the concept of the synchronous atom, an idea borrowed from particle accelerators. A related concept is ‘phase stability’, or the tendency of particles with a range of positions and velocities close to those of the synchronous atom to remain bunched with it during the operation of the machine. First discussed in the context of accelerators [218], models were later developed for Stark deceleration [32, 219] and then Zeeman deceleration [220]. The boundary of the phase-stable region is called the ‘separatrix’, the enclosed volume of which, or ‘acceptance’, is a common metric used in regard to molecular beams and traps. We can expand on the theme of decelerator loading by visualising the overlay of the phase-space coordinates of the molecular beam with the separatrix. This section describes a model for the longitudinal acceptance of the decelerator; extending the model to include the transverse dimensions will be explored in section 4.3.

The phase-stable region of a trapping potential is determined by the shape of the potential, and where analytic descriptions of the trapping field are not available the problem may be addressed with a Monte Carlo method. We will consider constant velocity guiding and deceleration of the low-field-seeking  $m_J = 2$  state of  $^3P_2$  metastable argon, but this method can be used for other low-field seeking species. For the appropriate decelerator parameters (driving voltage, peak current, initial and final travelling-wave velocity) the pulse-width modulation sequence to drive the power electronics is determined (section 2.3.6), from this the current generator code calculates the time-varying current through each coil phase and thus the magnetic field in the decelerator (section 3.6.1). For the time at which the moving trap of interest is fully formed the magnitude of the field along the  $z$ -axis is extracted. Figure 4.5 shows the force curves for this trap. For the deceleration case a constant pseud-

of force is applied, determined from the initial and final trap velocities and the length of the coils. As the trap forms the magnetic field profile is asymmetric, since it is at the beginning of the first coil (see figure 4.1) - when it reaches the centre of the active coil module the asymmetry is removed. The pseudoforce does not change the shape of the curve but slightly shifts the central equilibrium point by approximately 0.025 mm in the positive  $z$  direction, causing further asymmetry in the decelerated case. As the travelling-wave proceeds the trap depth and centre fluctuate (see section 4.2), and so the phase-stable regions determined here are only indicative of the conditions at the beginning of the sequence.

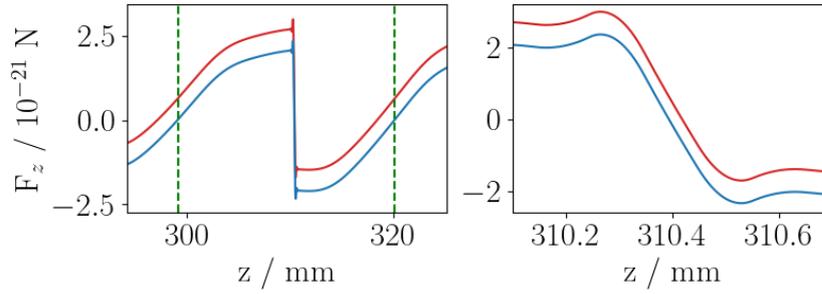


Figure 4.5: Left, force curve for argon atoms in the  $m_J = 2$  sublevel of the  $^3P_2$  term, within a single moving trap of the travelling wave, just as it is formed. The magnetic trap is moving with a constant velocity of  $336 \text{ m s}^{-1}$  (blue trace), and at  $337 \text{ m s}^{-1}$  with a deceleration of  $9497 \text{ m s}^{-2}$  (red trace). Green dashed lines indicate the maxima of the trap. Right, a magnified view of the centre of the left plot. Although the central equilibrium point of the trap is only shifted by 0.025 mm by the deceleration, the shape and area of the phase-stable region is much affected (*cf.* figure 4.8).

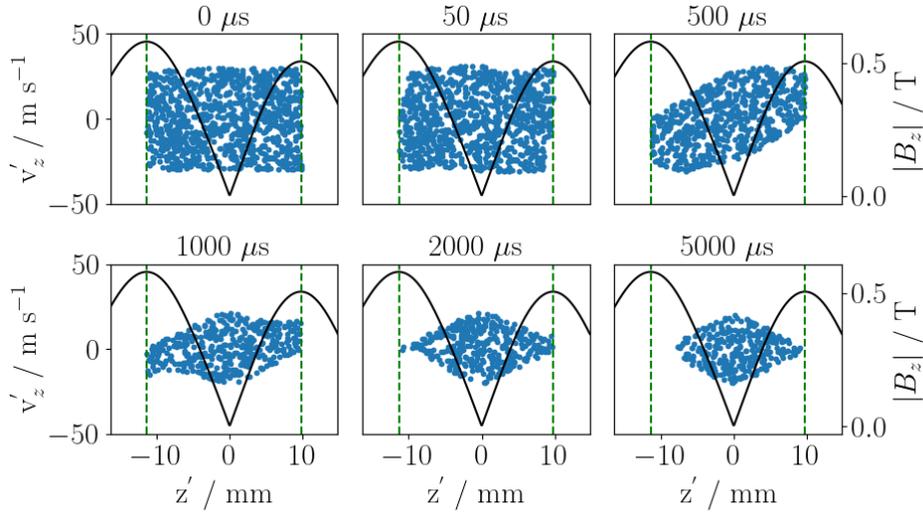


Figure 4.6: Monte Carlo determination of the approximate longitudinal acceptance for metastable argon, operating in constant-velocity guiding mode at  $336 \text{ m s}^{-1}$ , decelerator running at 600 V and 400 A. The longitudinal decelerator field (black trace) is calculated for the time at which the target trap is fully formed; this is populated with a uniform distribution of atoms in the low-field-seeking spin-stretched  $m_J = 2$  sublevel of the  $^3P_2$  term (blue). Their trajectories are numerically solved until enough time has passed that only phase-stable atoms remain; at every time step atoms that pass the trap maxima (dashed green) are discarded. The largest stable orbit approximates the separatrix.

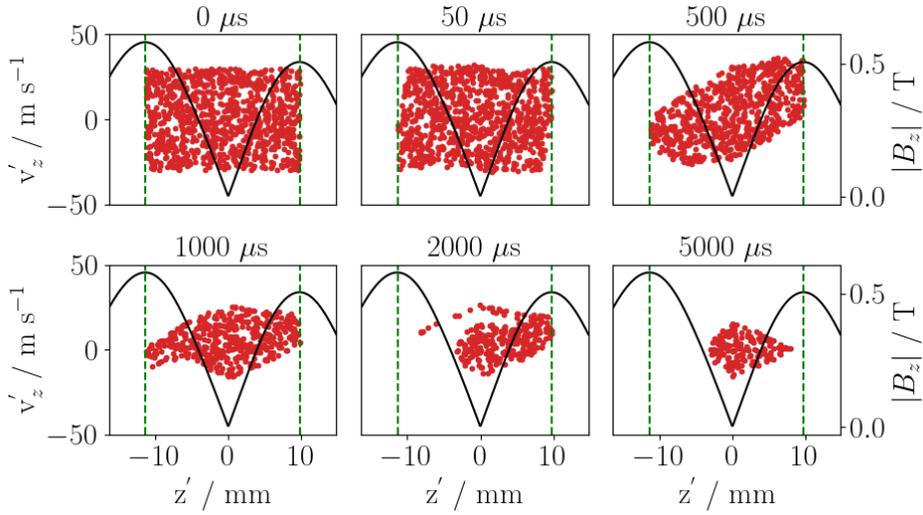


Figure 4.7: Monte Carlo determination of the approximate longitudinal acceptance for metastable argon, as per figure 4.6 but the decelerator is operating in deceleration mode from  $337 \text{ m s}^{-1}$  to  $324 \text{ m s}^{-1}$ , corresponding to an acceleration of  $-9497 \text{ m s}^{-2}$ . The magnetic trapping interaction is adjusted by adding a pseudo-force to match the deceleration of the wave; this is valid since the deceleration is shown to be approximately constant in a separate model that tracks the trap position over time. Again, the largest stable orbit approximates the separatrix.

Figure 4.6 shows the longitudinal decelerator field around the region of the trap just as it is fully formed, with minimum at  $z = 310 \text{ mm}$  from the source. The decelerator was programmed with a constant velocity guiding sequence with a target of  $335 \text{ m s}^{-1}$  and a maximum current of  $400 \text{ A}$  at  $600 \text{ V}$ . Section 4.2 discusses the determination of the actual trap dynamics by fitting the output of a full time-dependent simulation of the coil sequence; this method shows that the trap velocity for these parameters is actually  $336 \text{ m s}^{-1}$ . We proceed by treating the field as fixed and move into a reference frame  $z', v'_z$  centred on the trap, so that a particle moving with a velocity  $v'_z = 0$  is stationary with regard to the travelling wave. A uniform random sample of argon atoms is generated with  $z'$  coordinates between the trap maxima and velocities  $v'_z \pm 30 \text{ m s}^{-1}$ . The atoms are now allowed to evolve, the interaction with the field being solved as per section 3.4, with any atoms passing the maxima being discarded. The simulation is run until only phase-stable atoms remain; the area of the region they occupy approximates the longitudinal acceptance, and the largest stable orbit describes the separatrix. Note that the time required for all phase-unstable atoms to escape is much larger than a typical deceleration sequence. The asymmetry mentioned in the context of the force curves is visible in the plot of the field magnitude  $B_z$  and the shape of the separatrix. As the trap moves position further into the active coils the asymmetry will be reduced.

The decelerator was then configured to decelerate from  $335$  to  $325 \text{ m s}^{-1}$ . Simulation and fitting of the field sequence showed that the trap actually moves from  $337$  to  $324 \text{ m s}^{-1}$ , with a corresponding acceleration of  $-9497 \text{ m s}^{-2}$ . In the trap frame the deceleration manifests as a pseudoforce. We proceed as before and generate a sample of atoms then solve their motion, including the acceleration due to the pseudoforce. In figure 4.8 the trajectories of the largest stable orbits, moved back into the  $z, v_z$  frame, are plotted as the separatrices, showing that deceleration produces a smaller overall acceptance and much more pronounced asymmetry of the phase-stable region. The uniform initial atom distributions in figures 4.6 and 4.7 have two-dimensional phase-space volume of  $21 \text{ mm} \times$

$60 \text{ m s}^{-1} = 1260 \text{ mm m s}^{-1}$ . For the guiding and deceleration cases 34.4 % and 16.9 % of the atoms are stable in the trap; multiplying these fractions by the source volume gives estimated longitudinal acceptance values for spin-stretched argon of  $0.344 \times 1260 = 433 \text{ mm m s}^{-1}$  and  $0.169 \times 1260 = 213 \text{ mm m s}^{-1}$ .

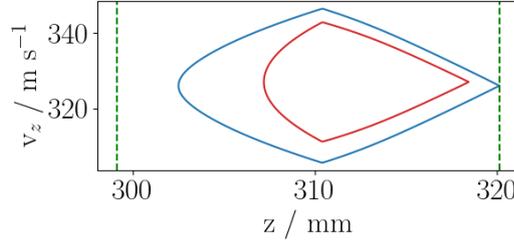


Figure 4.8: Comparison of longitudinal separatrices for the moving trap at the instant of formation for constant-velocity guiding mode at  $336 \text{ m s}^{-1}$  (blue trace) and deceleration mode from  $337$  to  $324 \text{ m s}^{-1}$  (red trace). Green dashed lines show trap maxima.

### 4.1.3 Longitudinal Decelerator Loading Model

In figure 4.9 the estimated guiding-mode separatrix is overlaid on a plot of a one-dimensional Monte Carlo simulation of argon guiding. With an electronics delay setting determined as per the method in section 4.1.1, the plot shows the moment at which the trap is fully formed. Although this one-dimensional model is only a limited representation of the real physical layout it is helpful in explaining some of the characteristics of our experiment. This is expanded on in section 3.10 which compares Monte Carlo models and experimental data. In the context of decelerator loading, we can use the approximate separatrix of the initial trap to study the effect of varying the power electronics delay on the fraction of the molecular beam that is available to be loaded, *i.e.* the relative number of atoms that are inside the phase-stable region as the trap is fully formed. Although it is simple to determine ‘by eye’ whether this is true for any given particle, achieving this computationally for many millions of simulated atoms is more difficult.

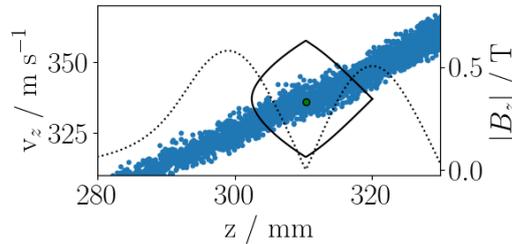


Figure 4.9: Guiding separatrix overlaid on a decelerator simulation. Blue dots are the phase-space coordinates of argon atoms, filled dot is the ‘ideal’ atom as in section 4.1.1, dashed line is longitudinal field. An efficient way to determine which particles lie within the closed surface of the separatrix was developed, see section 3.9.

Point-in-polygon was initially considered as a solution to this problem, but proved to be inefficient and difficult to implement. The method of point location via coordinate transformation (see section 3.9) was developed as an alternative. The method is illustrated in figure 4.10, which shows a separatrix and its transformation.

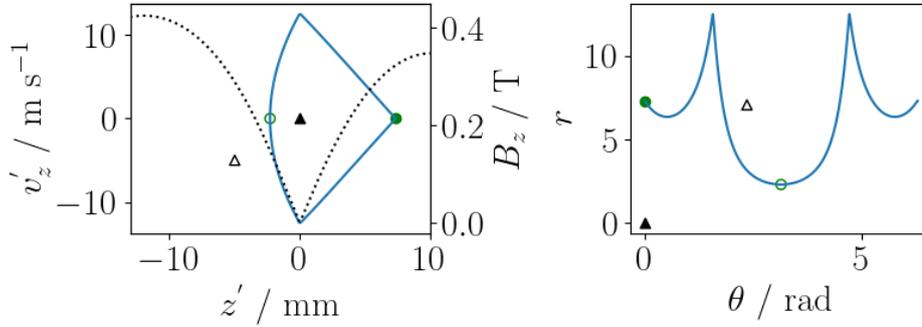


Figure 4.10: Axial separatrix for one well in the decelerator field. Left, the separatrix is shown in a Cartesian coordinate system, in a reference frame centred on the well. The axial field is shown as a dotted line. Right, the same separatrix plotted in a polar coordinate system. Filled and hollow dots mark equivalent points on the boundary of the phase-stable region, filled and hollow triangles mark the central point in the trap and an arbitrary location outside the phase-stable region. By loading the polar separatrix into a one-dimensional interpolator it is simple to determine whether a point lies above or below the phase-stable boundary.

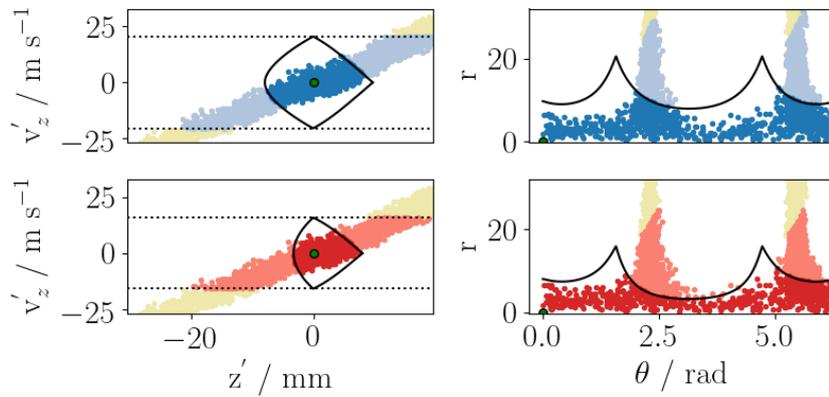


Figure 4.11: One-dimensional Monte Carlo simulation of  $^3\text{P}_2$  argon with central velocity of  $372 \text{ m s}^{-1}$  and temperature of  $6.1 \text{ K}$ , top, decelerator is running in guiding mode at  $336 \text{ m s}^{-1}$ , bottom, deceleration from  $337 \text{ m s}^{-1}$  to  $324 \text{ m s}^{-1}$ , both at  $600 \text{ V}$  and  $400 \text{ A}$ . Electronics delay  $t_d = 813 \mu\text{s}$  and separatrices for the  $m_J = 2$  sub-level estimated as per sections 4.1.1 and 4.1.2. Left, view in a reference frame centred on the moving trap. To optimise the loading of the trap the delay is adjusted to maximise the ratio of stable  $m_J = 2$  atoms (dark blue / dark red) to those that could potentially be loaded (between the dotted lines, all blue / red). Right, polar coordinate system, the separatrix is now a horizontal boundary and the phase-stable region lies below it. The yellow atoms are too slow / fast too to be phase-stable, for any value of  $t_d$ .

Figures 4.11 and 4.12 show the method applied to the one-dimensional Monte Carlo simulation from figure 4.9; only atoms with velocities in the region between the dotted lines, corresponding to the trap high and low points, can potentially be loaded, these are shown in blue / red for guiding / deceleration. If the total number of atoms between these bounds is  $N_p$  and the number within the

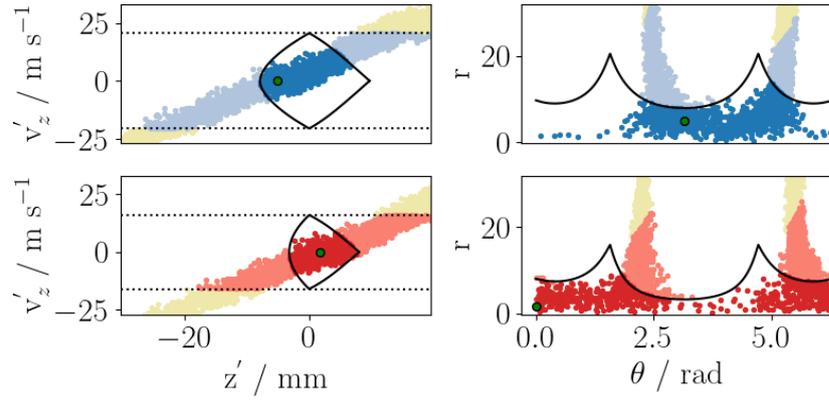


Figure 4.12: As per figure 4.11 but delay adjusted to optimise number of  $m_J = 2$  atoms within the phase-stable region. Top,  $t_d = 798 \mu\text{s}$  gives a 2 % increase. Bottom,  $t_d = 818 \mu\text{s}$  gives an almost negligible improvement, due to the shape of this separatrix. See figure 4.13.

separatrix (dark blue / red) is  $N_s$  then we wish to maximise the ratio  $N_I = N_s/N_p$ . This was done by varying the decelerator delay  $t_d$ , starting with the value  $t_i$  at which the ‘ideal’ atom arrives at the centre of the trap just as it forms (section 4.1.1). In 4.11  $t_d = 813 \mu\text{s}$  for both cases, in 4.12 this has been offset by  $-15 \mu\text{s}$  for guiding and  $5 \mu\text{s}$  for deceleration, giving an increase in capture of 2 % and 0.1 % respectively. Greater decelerations result in more asymmetry in the separatrix and a larger delay is needed to optimally load the trap. The variation of  $N_I$  with  $t_d$  is shown in figure 4.13, where  $\Delta t_d = t_d - t_i$ . These curves are dependent on the flow velocity, translational temperature and species of the beam, but the code can easily be adapted to suit other values of these parameters as required.

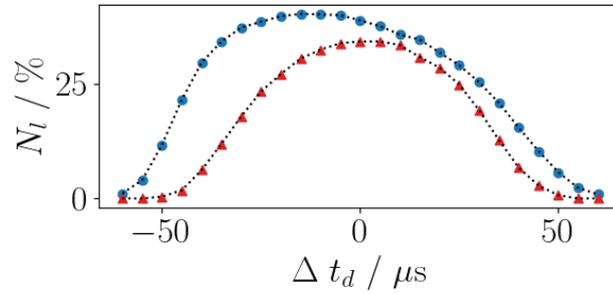


Figure 4.13: Dependence of the fraction of argon atoms loaded into the decelerator trap  $N_I$  with the power electronics delay, for constant-velocity guiding mode at  $336 \text{ m s}^{-1}$  (blue plot) and deceleration mode from  $337$  to  $324 \text{ m s}^{-1}$  (red plot).  $\Delta t_d = 0$  corresponds to a delay value that brings the ‘ideal’ or central atom in the beam into the centre of the trap at the time that it is fully formed, as per section 4.1.1.

## 4.2 PWM Effects

We utilise the technique of pulse-width modulation (PWM) to generate the required time-varying magnetic fields in our travelling-wave Zeeman decelerator. The desired ideal sinusoidal current waveform is calculated and a sequence of switching events for the power electronics is generated, as described in section 2.3.6. The degree to which the achieved waveform approximates the ideal sinusoid is ultimately determined by two factors; the electrical characteristics of the coils, and the response and recovery times of the switching components. This constrains the possible values of the frequency and peak current of the output waveform, and its deviation from the target. Figures 4.14 and 4.15 show two examples of PWM generation, with the ideal target waveform and estimated current compared to the actual current measured at the coils.

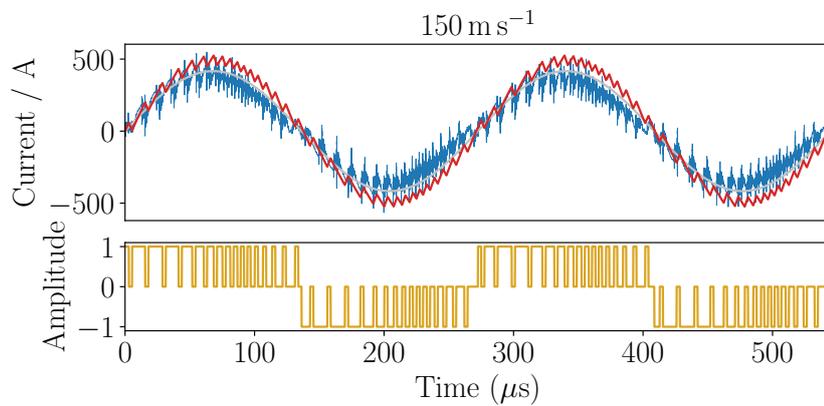


Figure 4.14: PWM example; target waveform (grey), the pulse sequences that are calculated to produce it (orange), estimated achieved current (red) and actual current measured at the decelerator coils (blue). The decelerator was operating in guiding mode at  $150 \text{ m s}^{-1}$ , 500 A peak at 650 V.

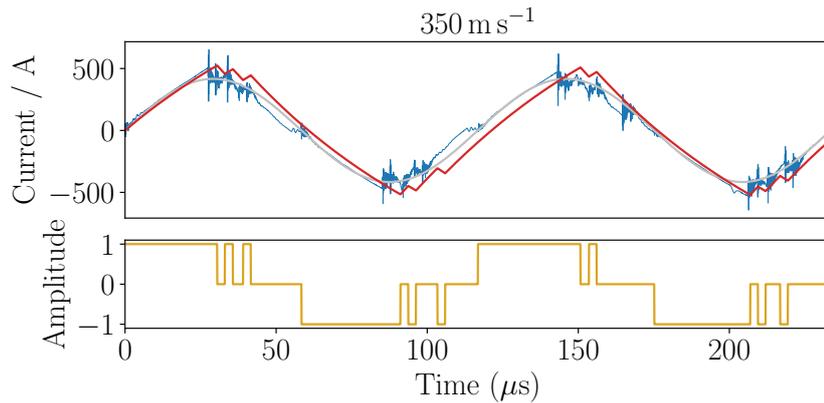


Figure 4.15: As per figure 4.14, but guiding at  $350 \text{ m s}^{-1}$ .

The effects of the PWM were previously studied by McArd [2] via simulation of the time evolution of one trap in the decelerator wave, comparing the synthesised waveform with a ‘pure’ sinusoid. This was performed for two different wave velocities, allowing the fluctuations of the trap centre and depth to be determined. Due to limitations of the codes available at the time these simulations only covered a section of the decelerator. It was concluded that as the trap velocity increased the

frequency of the jitter in trap position decreased, with the magnitude of the oscillations increasing slightly. The trap depths were on average lower for the synthesised compared to pure waveforms, with a similar inverse relationship between velocity and frequency.

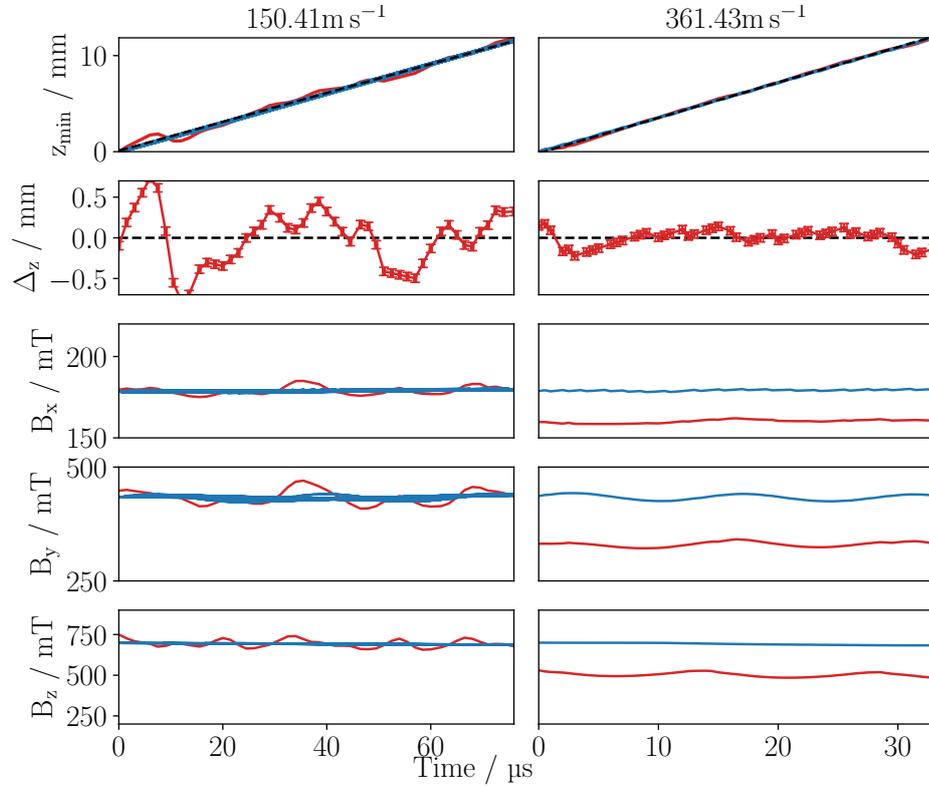


Figure 4.16: Effects of the PWM technique on one trap in the decelerator wave as it passes through one quarter of a period of one coil. For all plots the PWM-generated waveform data are in red, the reference pure sinusoid data are in blue. First row, the  $z$ -position of the trap minimum over time, the dashed line is fitted to the PWM data to give the mean trap velocity. Second row, the variation between the PWM data and the fitted mean. Third to fifth rows show the fluctuation of the trap depth in the  $x$ -,  $y$ - and  $z$ -directions. Left / right, target velocity of 150 / 350  $\text{m s}^{-1}$ , fitted velocities shown at the top. In this simulation the wire quadrupole is active at 700 A.

Utilising our fast interpolator code and the methods described in sections 2.3.6 and 3.6.1 we have been able to extend this analysis to the entire decelerator sequence, with a higher sampling frequency and grid resolution. We begin by plotting the decelerator field along the beam propagation (that is,  $z$ ) axis for several tens of  $\mu\text{s}$  to determine the time at which a target trap forms and its longitudinal position. The trap chosen is the second one to form, due to its favourable combination with the wire quadrupole field (see section 2.3.7). The code then runs a full sequence of the wave propagation, and at each timestep the longitudinal field is examined to determine the position of the trap minimum. Also recorded is the magnitude of the field at the front of the trap, and at the transverse positions corresponding to the radius of the sample delivery tube in the  $x$  and  $y$  directions. For guiding mode the  $z$  versus time data can be fitted to determine the average trap velocity, or for deceleration mode they can be fitted for the initial / final velocities and deceleration.

Figure 4.16 shows some results across one quarter of a period of one coil, chosen to be comparable to the figures shown in section 4.4 of McArd, for two sequences corresponding to target wave velocities of  $150$  and  $350 \text{ m s}^{-1}$ . The wire quadrupole is active at  $700 \text{ A}$ . The centres of the first traps in the decelerator wave are plotted against time in the upper row, and fitted to give the mean velocity. The variation of the trap centres from the positions expected from the fitted mean are shown in the second row, and we see a difference from the earlier simulations; at higher velocities the oscillations do appear to have a lower frequency, but they also show a significantly smaller amplitude. The error bars represent the spatial grid across which the simulation was performed, which is  $0.1 \text{ mm}$  in the  $z$ -direction, and the timestep, which was  $0.5 \mu\text{s}$ . Rows three through five show the trap depths along the Cartesian axes for the synthesised (red) versus pure (blue) waveforms. McArd reported that in the PWM case the trap became shallower, with the effect increasing at higher velocities; we find the same, but the effect is much more pronounced than previously estimated. As before, the frequency of these oscillations is inversely related to the velocity.

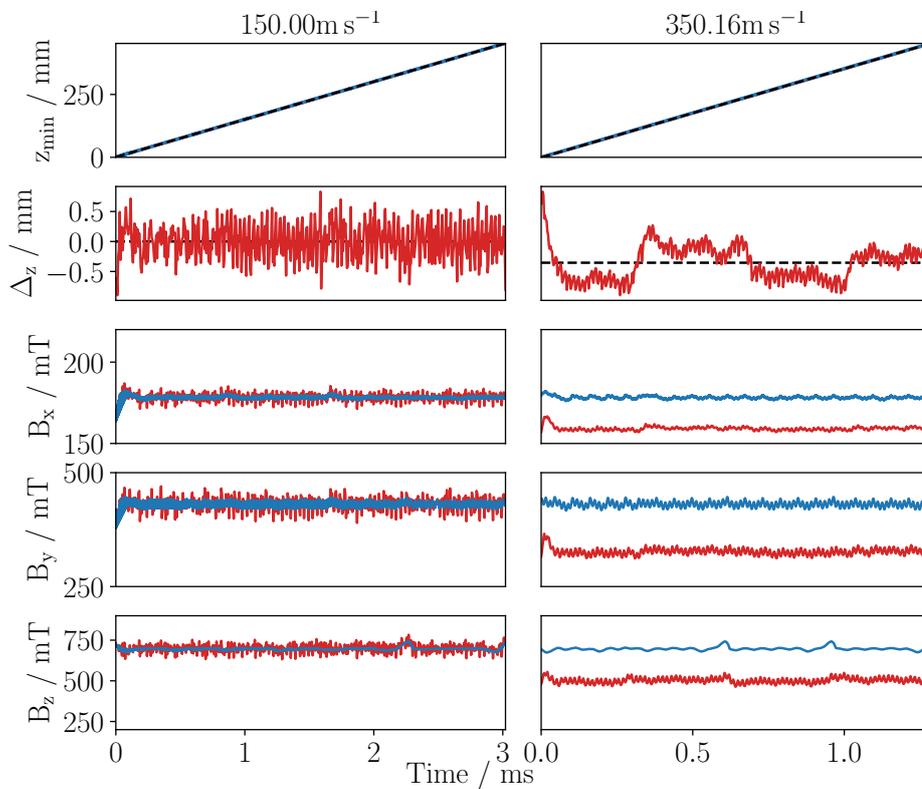


Figure 4.17: As per figure 4.16, but showing the motion of one trap in the wave across the entire decelerator.

Figure 4.17 shows the same analysis as figure 4.16 but extended across the entire four-module decelerator. The achieved trap velocity is closer to the target than in the case of the smaller sample, indicating that the variation of the trap minimum is not uniform throughout an entire coil. The difference between the desired and achieved velocities also gets larger with the target velocity. The boundaries between coil modules can be seen in the peaks in longitudinal depth for the faster wave.

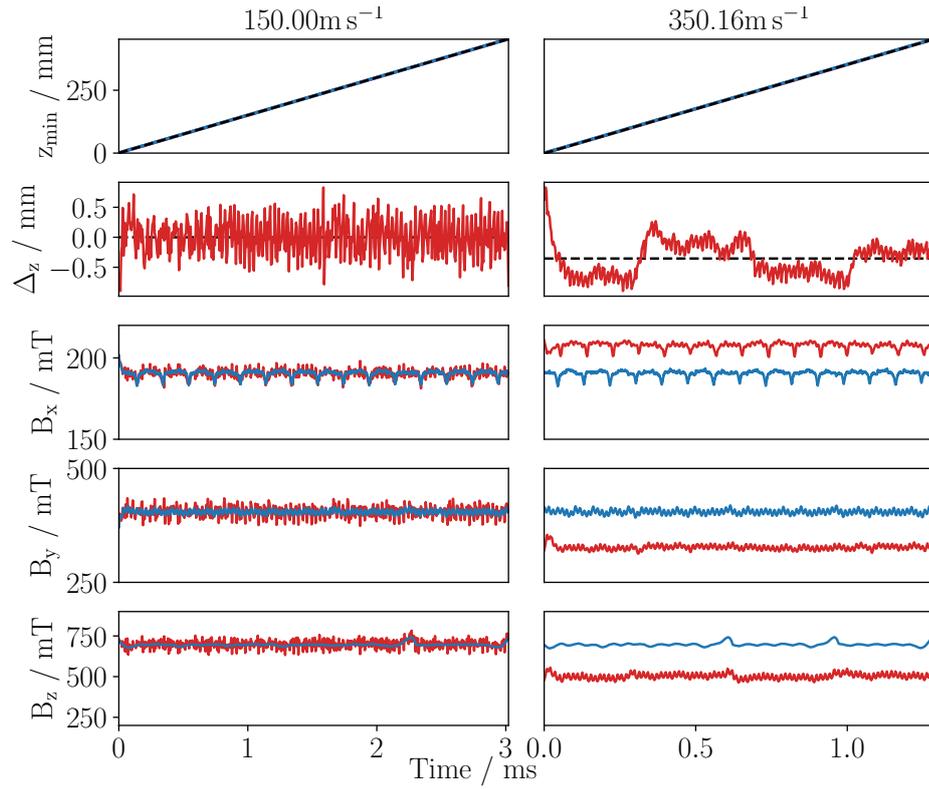


Figure 4.18: As per figure 4.17 but with permanent magnet rather than wire quadrupole (see section 2.3.7).

Figure 4.18 shows the effect of replacing the wire quadrupole with the permanent magnet quadrupole (described in section 2.3.7). The achieved wave velocities and longitudinal trap depths are practically the same. Notable is the much improved trap depth in  $x$ . In addition to the periodic variation along  $z$  due to the decelerator coil boundaries a regular variation in transverse depth can be seen, indicating the regions where the quadrupole magnets adjoin.

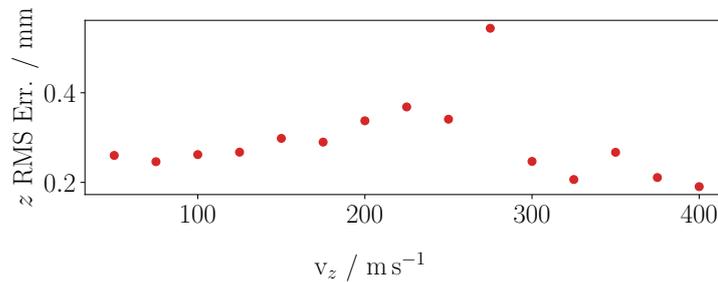


Figure 4.19: Estimated root-mean-square error in the central position for one trap in the travelling wave, compared to the expected position from the fitted mean velocity. All simulations have the decelerator electronics set to operate with a peak current of 500 A at 650 V, PWM tolerance at 8 %.

The simulations were repeated for a range of travelling-wave velocities. Figure 4.19 shows the RMS error in the longitudinal position of the trap, and it can be seen that the positional jitter increases with wave velocity from  $\approx 150 \text{ m s}^{-1}$  to a surprisingly high peak at  $275 \text{ m s}^{-1}$  before decreasing again. Figure 4.20 shows the mean and standard deviation of the oscillations in trap depth, which display

similar behaviour in all three dimensions. The mean depth steadily decreases as the wave velocity increases above  $200 \text{ m s}^{-1}$ , as does the standard deviation except for an unexpected outlier, again at  $275 \text{ m s}^{-1}$ .

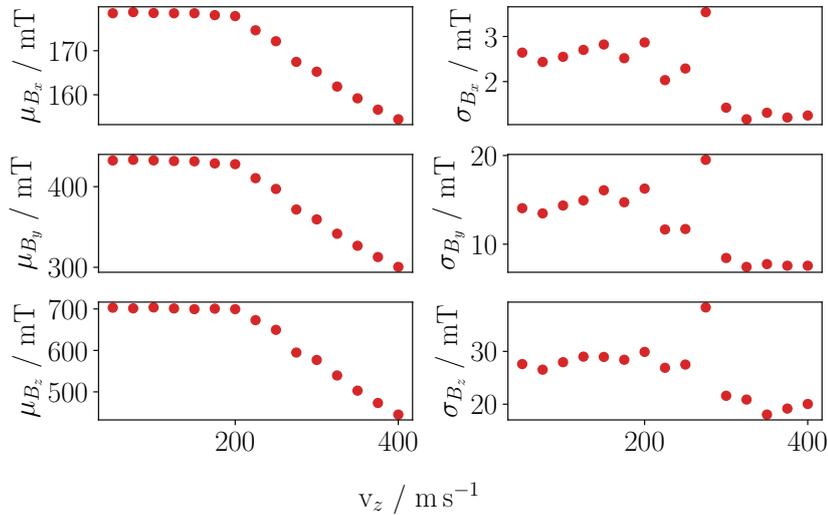


Figure 4.20: Estimated variation of the depth of one trap in the decelerator wave, as a function of velocity. Shown are the mean and standard deviation of the depth along the three coordinate axes. Peak current of 500 A at 650 V, PWM tolerance at 8 %.

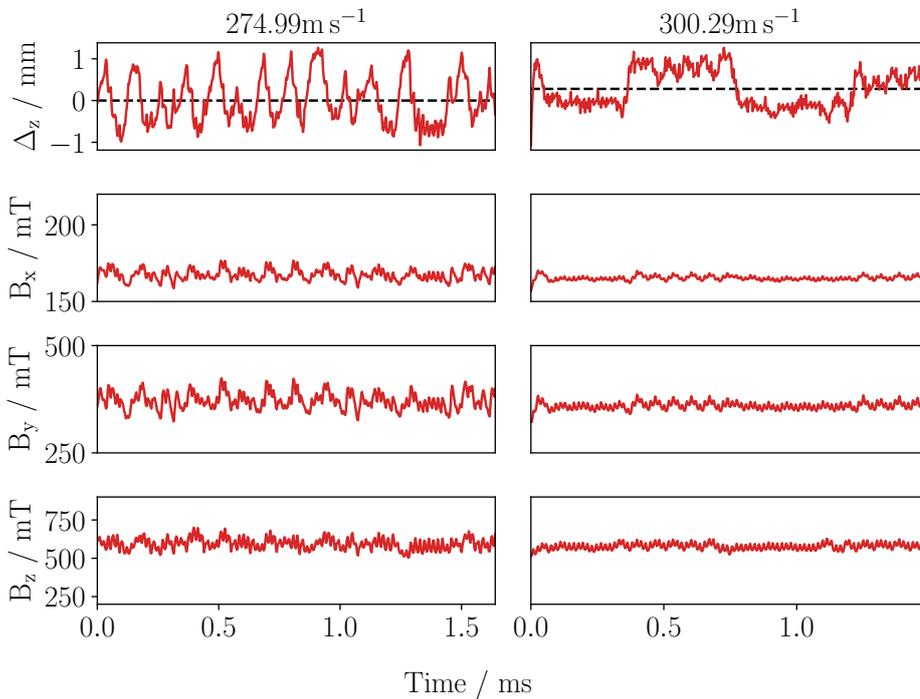


Figure 4.21: Analysis of variation of trap position and depth, guiding mode at  $275 / 300 \text{ m s}^{-1}$ , 500 A at 650 V.

Figure 4.21 compares the trap properties for guiding at  $275$  and  $300 \text{ m s}^{-1}$ . Larger amplitude oscillations in depth and position can clearly be seen for the slower wave. An explanation can be found by considering the PWM pulse sequence calculated for  $275 \text{ m s}^{-1}$ , shown in figures 4.22 and 4.23.

Using the typical tolerance envelope value of 8 % of the peak current, the code is forced to add extra short pulses in order to try and match the reference waveform. Increasing the tolerance value allows a synthesised waveform to be built from fewer pulses, meaning that it is more smooth. This illustrates the need to inspect the output of the PWM code for a set of target experimental parameters.

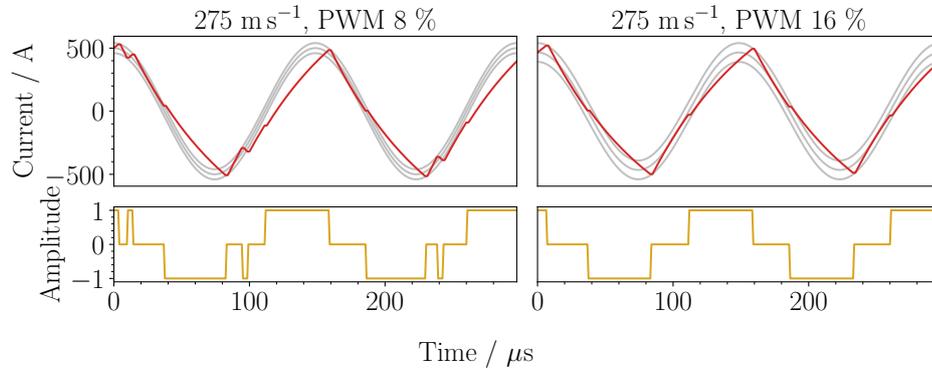


Figure 4.22: PWM current synthesis, guiding mode velocity of  $275 \text{ m s}^{-1}$ ,  $500 \text{ A}$  at  $650 \text{ V}$ ; target waveform (grey), pulse sequences (orange), estimated current (red). Left / right, PWM tolerance at 8 % / 16 %.

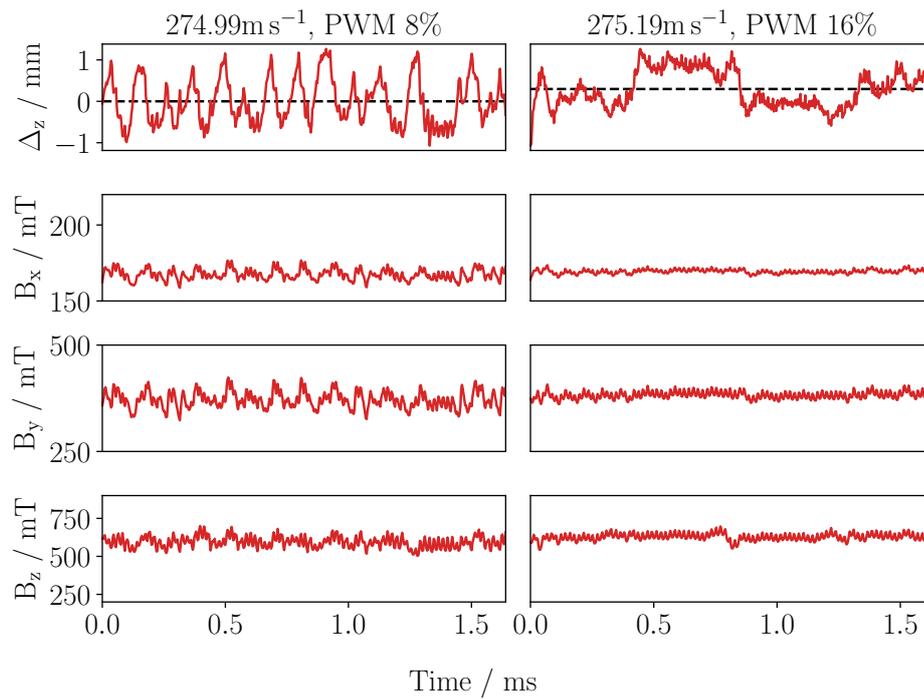


Figure 4.23: Analysis of variation of trap position and depth, guiding mode at  $275 \text{ m s}^{-1}$ ,  $500 \text{ A}$  at  $650 \text{ V}$ , left / right PWM tolerance set to 8 % / 16 %.

In summary; the PWM method causes the longitudinal position of the moving trap to oscillate with an amplitude somewhat proportional to the wave velocity. The trap depth also decreases as the wave gets faster. Inspection of the calculated pulse solution for a given set of parameters may indicate the problem; increasing the tolerance to which the target waveform is approached may give a smoother current profile, with reduced oscillation.

### 4.3 Static Acceptance Model

Section 4.1.2 described a model for longitudinal acceptance which is not very meaningful except as a tool for optimising the decelerator loading process. This section describes a simple model to estimate the acceptance across all dimensions for a single trap in the decelerator wave.

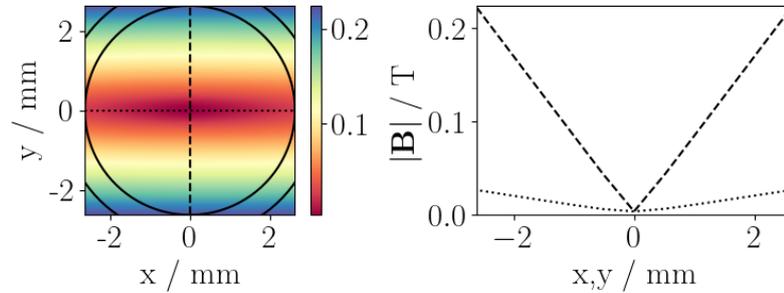


Figure 4.24: Magnetic field magnitude at the centre of the moving trap, decelerator guiding mode at  $336 \text{ m s}^{-1}$  with no additional quadrupole; at the instant of formation the trap is centred at  $z = 310 \text{ mm}$ , see figure 4.1. Left, the  $xy$ -plane, the solid circles represent the sample delivery tube (inner diameter of  $5.2 \text{ mm}$ ) that runs the length of the decelerator. Right, field strengths along the  $x/y$  (dotted / dashed) axes; the decelerator geometry gives rise to a trapping field that is much more strongly confining in the  $y$  direction. See sections 2.3.4 and 2.3.7.

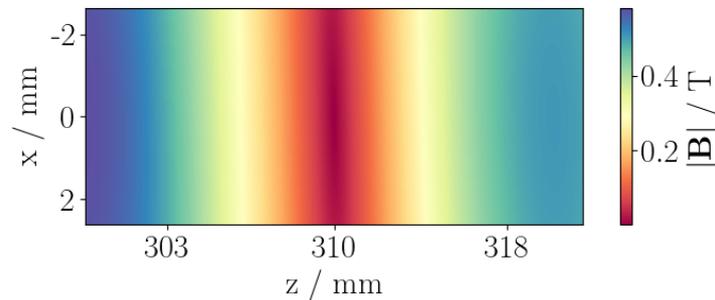


Figure 4.25: Magnetic field magnitude of the moving trap in the  $xz$ -plane, the plot is bounded in  $z$  by the trap longitudinal maxima, and in  $x$  by the inner diameter of the sample delivery tube. It can be seen that the trap is shallow in  $x$ , in contrast to  $y$  (figs 4.24, 4.26).

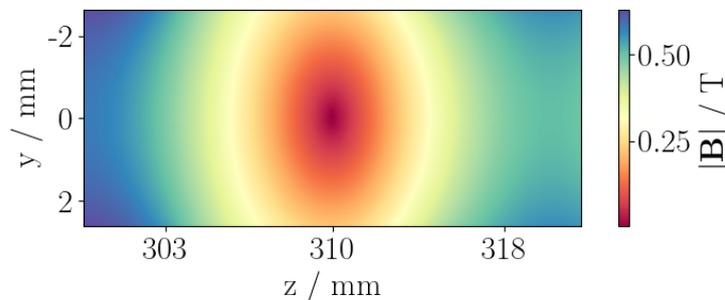


Figure 4.26: Magnetic field magnitude of the moving trap in the  $yz$ -plane, the plot is bounded in  $z$  by the trap longitudinal maxima, and in  $y$  by the inner diameter of the sample delivery tube. It is seen that the trap is much more strongly focussing in  $y$  than in  $x$  (figs 4.24, 4.25).

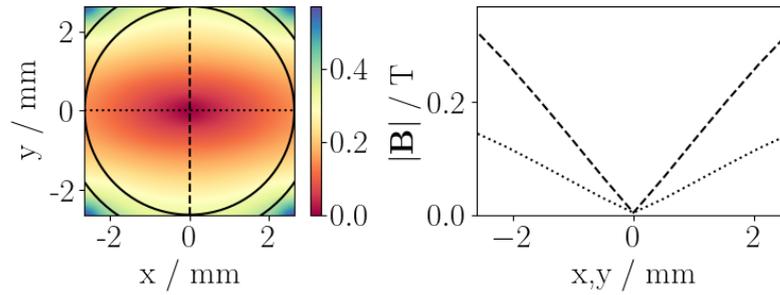


Figure 4.27: As per figure 4.24 but with wire quadrupole set to 700 A.

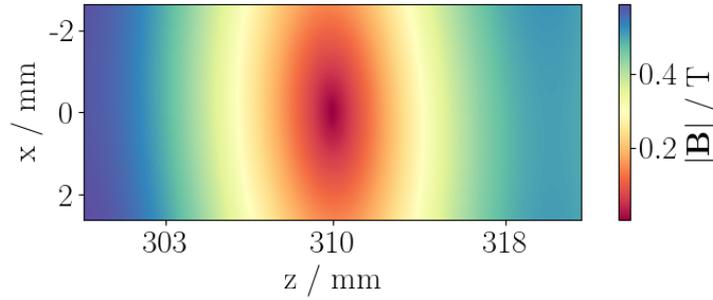
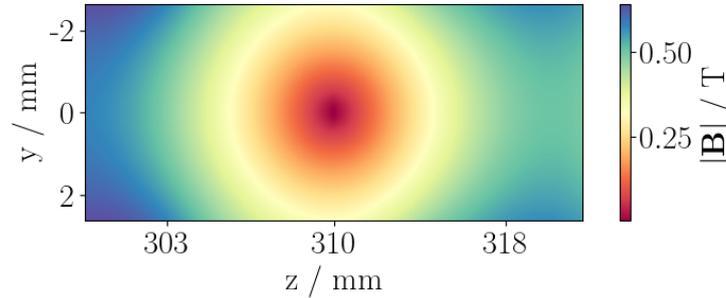
Figure 4.28: As per figure 4.25 but the wire quadrupole provides much stronger confinement in  $x$ .

Figure 4.29: As per figure 4.26 but with wire quadrupole active.

The profiles of the decelerator and quadrupole fields was covered in chapter 2 but figures 4.24 to 4.29 give a detailed view of a single trap. In the one-dimensional model we calculated the decelerator field along the beam axis at the moment the target trap was fully formed. Here, we instead extract the full three-dimensional scalar field  $B$  in the volume surrounding the trap. A random sample of atoms is then generated and allowed to evolve using the methods described in sections 3.4 and 3.8.2 until all ‘phase-unstable’ atoms have escaped. The trap field was modelled both with and without the wire quadrupole included. A similar study was performed by McArd [2], but our improved interpolator code allowed a greater range of configurations to be tested.

For convenience we move to a reference frame with coordinates  $x', y', z'$  centred on the trap, so that the trap centre is the origin along the beam axis  $z$  and a longitudinal velocity  $v'_z$  is stationary with respect to the trap (the transverse components are the same in both frames). The sample is created with initial positions and velocities chosen from uniform random ranges  $-6.5 \text{ mm} \leq [x', y', z'] \leq 6.5 \text{ mm}$  and  $-18 \text{ m s}^{-1} \leq [v'_x, v'_y, v'_z] \leq 18 \text{ m s}^{-1}$ . These values do not have to be chosen very

carefully as long as they are larger than the actual phase-stable region, with the caveat that the larger the initial volume in phase-space the more particles need to be added to attain meaningful statistics, increasing the processing time and memory requirements at the beginning of the simulation. Once running, however, non-phase-stable particles rapidly leave the trap region, being counted as lost once they reach the boundary of the modelled field region. The code runs until the number of particles remaining has stabilised; the acceptance is then estimated by multiplying the initial phase-space volume by the fraction of the particles that are still in the trap. For the case of deceleration rather than guiding a pseudoforce is added when solving the trajectories at each time step. A constant deceleration is assumed, the magnitude of which is determined by considering the initial and final trap velocities and the distance between the first and last decelerator coils. Figure 4.30 shows examples of the initial and final phase-space distributions for argon, for both guiding and deceleration. The shape of the final  $z, v_z$  components can be compared to the separatrices, discussed in sections 4.1.2 and 4.1.3.

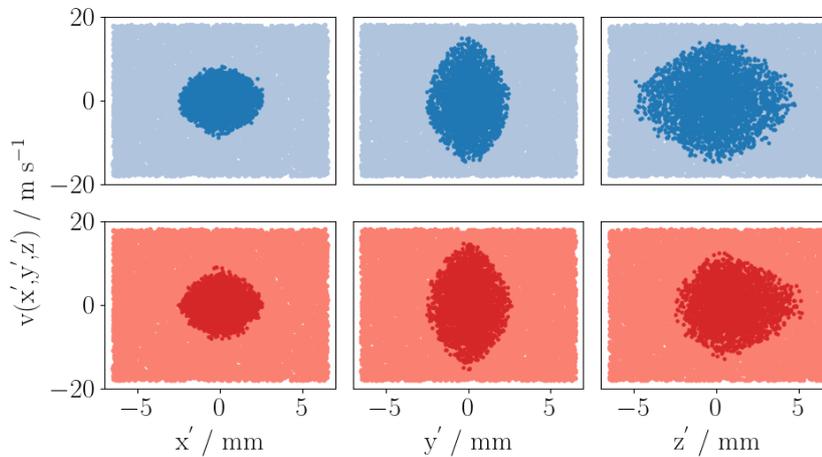


Figure 4.30: Estimation of decelerator phase-space acceptance for  $^3P_2$  argon by solving particle trajectories in a single trap of the decelerator plus wire quadrupole field, for guiding at  $336 \text{ m s}^{-1}$  (blue) and deceleration from  $337 \text{ m s}^{-1}$  to  $324 \text{ m s}^{-1}$  (red). An initial sample of atoms with a uniform spread of positions and velocities is created (light blue/red) and allowed to evolve in the trap until only phase-stable atoms remain (dark blue/red). The fraction that remains is multiplied by the initial phase-space volume to give an estimate of the acceptance.

Simulations were carried out for argon in both the existing 4-module decelerator and the proposed 9-module extended design, with a range of target final velocities. The initial trap velocity was set at  $350 \text{ m s}^{-1}$ ; our supersonic expansion source produces argon beams with approximately this flow velocity. Figure 4.31 shows the results, both with and without the wire or permanent magnet quadrupole guide active. It can be seen that for this relatively heavy species the acceptance of the short decelerator decreases rapidly with increased deceleration. The effect of the transverse focussing field is also evident. These results are close to those found by McArd [2]. Finally, for comparison figure 4.32 shows the results of applying this method to ground-state hydrogen (specifically the spin-stretched  $m_F = 1$  sub-level). As expected, both the spatial and velocity components of the acceptance of this light atom are significantly larger than those of argon. A more complete approach to determining the decelerator acceptance is described in section 4.4.

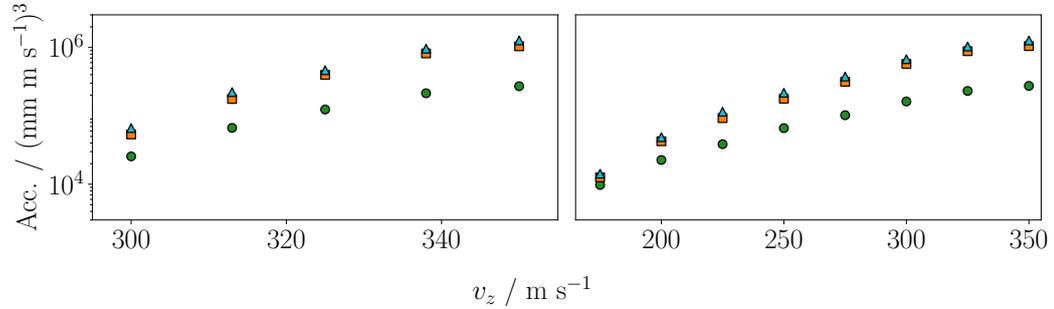


Figure 4.31: Six-dimensional phase-space acceptance of the travelling-wave decelerator for the  $m_J = 2$  state of  ${}^3P_2$  argon, estimated via simulation of one trap in which the field is static. The horizontal axis shows the final wave velocity, which has an initial value of  $350 \text{ m s}^{-1}$ . Left, the existing 4-module machine, right, proposed extended 9-module design. The field corresponds to a peak current of  $500 \text{ A}$  at  $650 \text{ V}$  in the deceleration coils. The squares / triangles / circles represent the field with at  $700 \text{ A}$  wire / permanent magnet / no additional quadrupole.

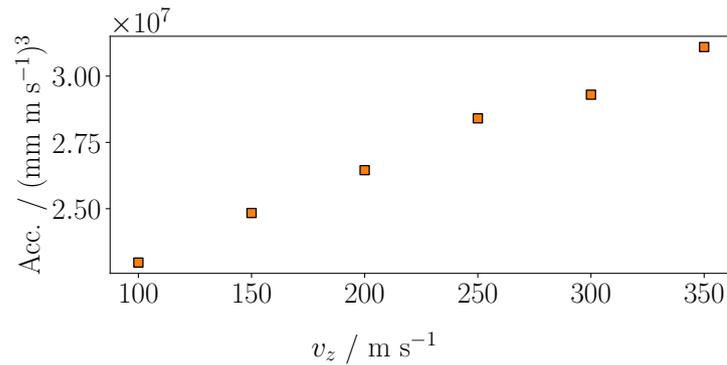


Figure 4.32: Six-dimensional phase-space acceptance of the 4-module travelling-wave decelerator for the  $m_F = 1$  state of  ${}^2S_{1/2}$  hydrogen, parameters otherwise as per those in figure 4.31 with wire quadrupole. As expected, for this lighter species the decelerator has significantly greater acceptance.

## 4.4 Dynamic Acceptance Model

Sections 4.1.2 and 4.3 described our attempts to estimate the phase-space acceptance of the travelling-wave Zeeman decelerator via static simulations. These involved taking a snapshot of the magnetic field around a single trap in the decelerator field and evolving the trajectories of a uniform distribution of particles, to see which ones would remain trapped. This method has the advantage of being relatively simple to set up, since the magnetic fields are only calculated once and treated as constant throughout the simulation, and this also means that the code has relatively low requirements for processor and memory, with associated short running times. In this idealised view of the decelerator the acceptance in guiding mode is independent of trap velocity, and is only a function of the trap depth (*i.e.*, the peak current in the coils). In deceleration mode the only relevant factors are the trap depth and the rate of deceleration. In reality, these static simulations do not capture the fluctuations in the depth and central position of the trap that are a side-effect of the pulse-width modulation (PWM) technique used to generate the time-varying currents in the decelerator coils, as shown in section 4.2. McArd [2] used similar static methods in earlier work on the decelerator, both to estimate the acceptance and to simulate the deceleration process, for which he produced a three-dimensional trap field that was the time-average of these fluctuations. A more desirable solution for determining the acceptance would therefore be to produce a fully dynamic time-dependent simulation of the deceleration (or guiding) process of a uniform sample of atoms. This is feasible using the methods outlined in chapter 3, and will be the subject of this section.

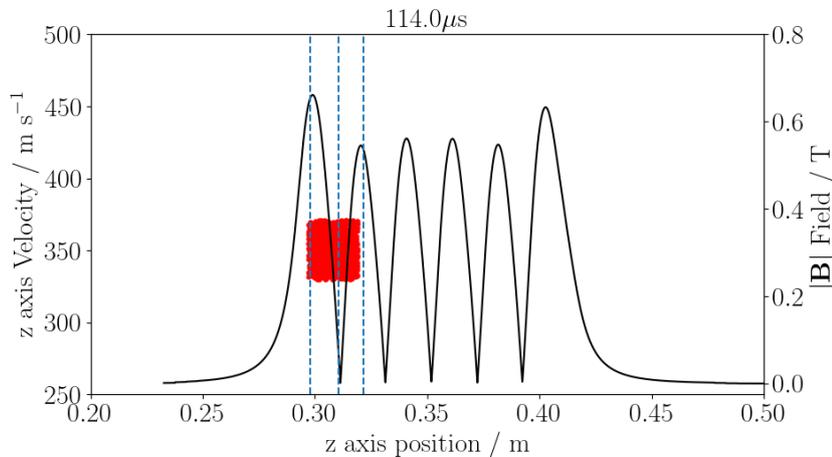


Figure 4.33: Snapshot of the Monte Carlo acceptance code, based on the decelerator dynamic simulation code. A sample of atoms with positions and velocities equally spaced in all dimensions is generated, centred on the target trap at the instant at which it is fully formed. The decelerator fields are calculated at each timestep and the motion of the trap centre and forward and rear maxima (dashed lines) are tracked. Particle motion including interaction with the fields is solved, and all particles that pass the maxima or move to a radial distance equating to the radius of the sample delivery tube are discarded. The simulation terminates when the field shuts off, and the fraction of particles remaining is used to calculate the acceptance.

The operation of the code is illustrated in figure 4.33. For a given set of decelerator parameters including initial and final wave velocity, voltage and peak current, the profile of the field as it rises is plotted. The time at which the target trap is fully formed and the position of its centre are determined by inspection and used as inputs to the acceptance code. At the start time a sample of atoms is created

with positions uniformly spaced along ranges centred on the beam axis in  $x$  and  $y$  and at the trap centre in  $z$ . The initial velocities are likewise generated, centred on zero in the transverse directions and with a mean value equal to the trap velocity at the beginning of the sequence. The longitudinal position and velocity components of the initial sample are shown in red on the plot. The on-axis field magnitude is shown, with the centre and forward and rear maxima of the trap marked by the dashed lines. At each time step the field is calculated and the code finds the new position of the trap centre by finding the nearest minimum to the last position along  $z$ . The forward and rear peaks are found by locating the nearest maxima in front of and behind the trap centre, and any particles that pass beyond these lines are removed. As is the case in the decelerator Monte Carlo code, at each timestep all particles that drift laterally to a distance corresponding to the sample delivery tube are also counted as lost and removed. This process continues until the field shuts down. The fraction of particles that remain is then combined with the phase-space volume of the initial sample to calculate the acceptance.

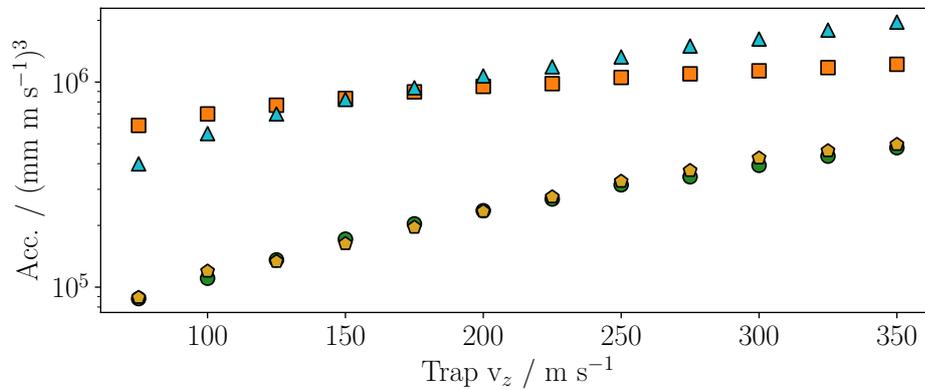


Figure 4.34: Estimated acceptance of the 4-module decelerator in constant-velocity guiding mode, for the  $m_J = 2$  sub-level of  $^3\text{P}_2$  argon. The circles / pentagons show the decelerator operating in PWM / ideal sinusoid mode with no quadrupole. The squares / triangles are for PWM mode with wire / permanent magnet quadrupole. The higher acceptance obtained with transverse focussing present is as predicted. Less expected is the performance of the permanent quadrupole dropping below that of the wire at low velocities.

Results of constant-velocity guiding of argon are shown in figure 4.34. The circles represent the decelerator field alone in PWM mode with no extra quadrupole. The pentagons also have no extra transverse focussing but the currents giving rise to the field are the ideal or ‘pure’ reference sinusoids used to create the PWM sequence. Section 4.2 described the fluctuations in the decelerator field arising from PWM, but the depths and positions of the trapping fields calculated using the pure waveform are not entirely free of variations in time, although on a much smaller scale. Nonetheless, it can be seen that the pure waveform gives rise to slightly larger acceptance at higher velocities when compared to PWM. The squares and triangles are both PWM simulations, and have, respectively, the wire and permanent magnet quadrupoles included. The greater acceptance produced by the permanent magnet system at higher velocities is as expected, given the stronger focussing fields it produces when compared to the wire one. More surprising is the drop in the performance of the former when compared to the latter at lower velocities. The cause is not immediately apparent; referring to figures 2.15, 2.16 and 2.17 in section 2.3.7 we would expect the permanent quadrupole to give better confinement in all cases.

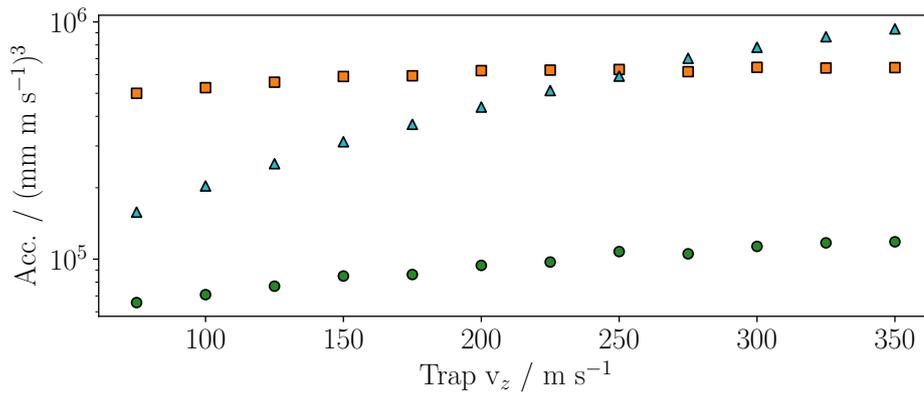


Figure 4.35: Estimated acceptance for argon in guiding mode, as per figure 4.34 but for the proposed extended 9-module decelerator. The circles, squares and triangles represent no, wire and permanent quadrupoles, respectively. The weaker performance of the permanent quadrupole at slower beam velocities, as seen in the 4-module case, is more prominent.

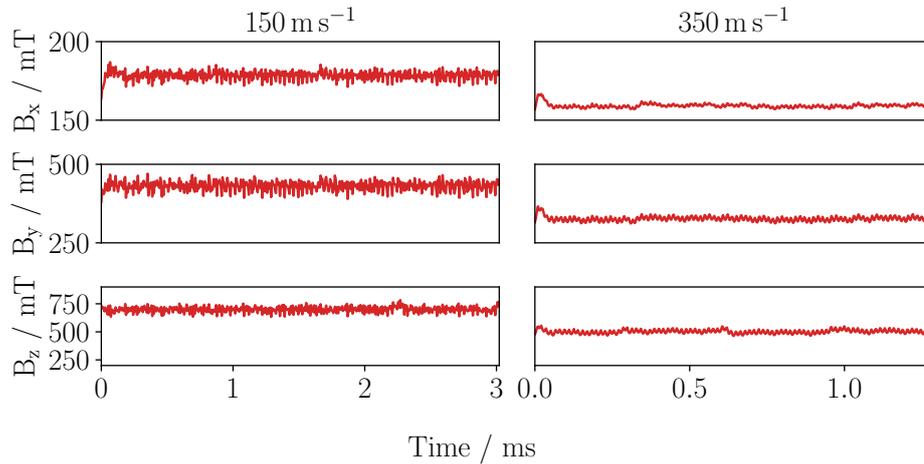


Figure 4.36: Effects of PWM on trap depth for a single trap in the decelerator wave, with the wire quadrupole applied. The average depth in  $x$  and  $y$  is higher than without the quadrupole, the field in  $z$  is unaffected.

In section 4.3 the change in quadrupole performance appeared to be linked to the magnitude of the current in the decelerator coils. Another possibility for the declining performance of the permanent as compared to wire quadrupole at low beam speeds in the dynamic simulation is the fact that there is a spatially periodic fluctuation in the shape of the transverse field arising at the boundaries between magnet sections. The wire system is reasonably uniform along its length. Possibly this oscillation contributes to a lowered acceptance for slower beams, in which the atoms spend more time in this region. This is illustrated in figures 4.36 and 4.37 which use the same analysis as in section 4.2. As expected, there is a negligible effect on the depth of the trap in the  $z$ -direction since neither quadrupole has a component in  $z$ . Compared to figure 4.17 the trap depths in  $x$  and  $y$  are on average larger for both quadrupoles than in the case with neither applied, but the periodic variation in  $x$  can be seen for the permanent case. Given the smallness of the dip in trap depth it seems unlikely that this causes significant loss of focussing but it cannot be ruled out without a more careful analysis of trap loss mechanisms.

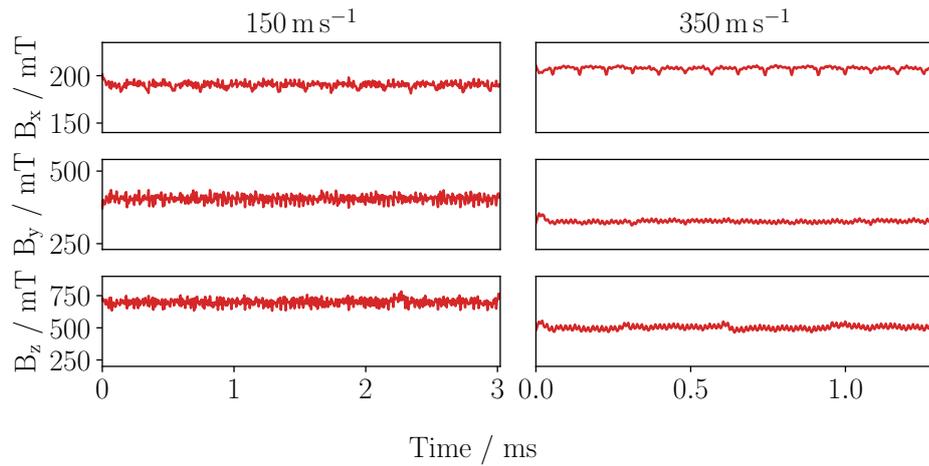


Figure 4.37: Effects of PWM on trap depth for a single trap in the decelerator wave, with the permanent quadrupole applied. The average depth in  $x$  and  $y$  is higher than without the quadrupole, the field in  $z$  is unaffected. There is a periodic oscillation of the field in the  $x$ -direction arising from the boundaries between magnets, and this may contribute to the worsening performance of the permanent quadrupole at lower beam speeds, since particles spend longer in these regions.

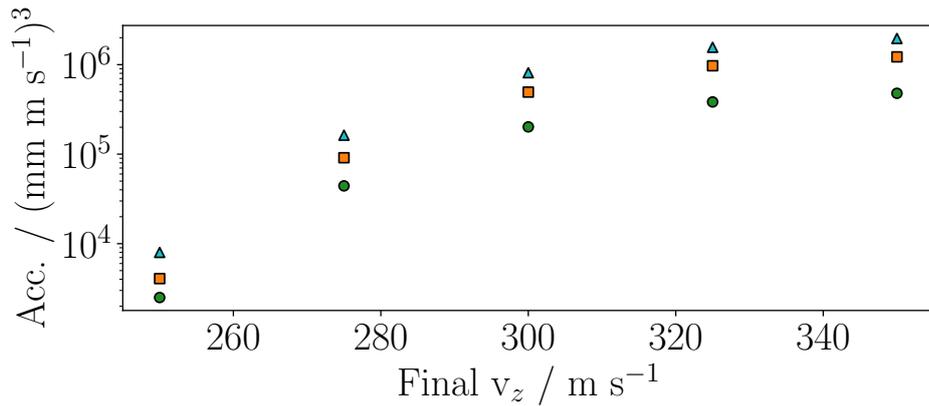


Figure 4.38: Estimated acceptance for argon in deceleration mode. The circles show the decelerator operating with no quadrupole. The squares / triangles have the wire / permanent magnet quadrupole. The discrepancy between quadrupoles seen in the guiding case does not appear for deceleration, however the acceptance in general drops very rapidly for larger decelerations for this combination of relatively heavy species and short decelerator.

Figures 4.38 and 4.39 show the results of deceleration of argon in the four- and nine-module decelerator, with or without the wire or permanent quadrupoles. As expected, the acceptance falls with the target trap velocity, with the longer design giving better results due to its gentler deceleration. The data suggest that the same effect that causes the permanent quadrupole to perform worse than the wire one at lower beam velocities is still present, but this is less significant for several reasons. The traps in the wave spend less time at the lower velocity, starting out faster, so whatever weakness in the potential the particles are exploiting, they have less time to do so. Also, the acceptance plummets by several orders of magnitude for significant decelerations of such a relatively heavy species as argon, since the deceleration effects dominate other trap loss mechanisms, and so the very low-velocity effects are irrelevant.

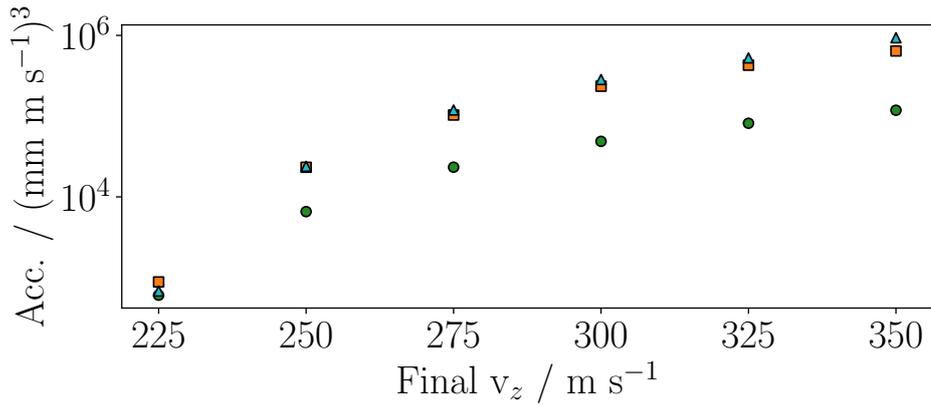


Figure 4.39: Estimated acceptance for argon in deceleration mode, as per figure 4.38 but for the proposed extended 9-module decelerator. The circles, squares and triangles represent no, wire and permanent quadrupoles, respectively. The fall in acceptance for larger decelerations is less pronounced for the longer machine.

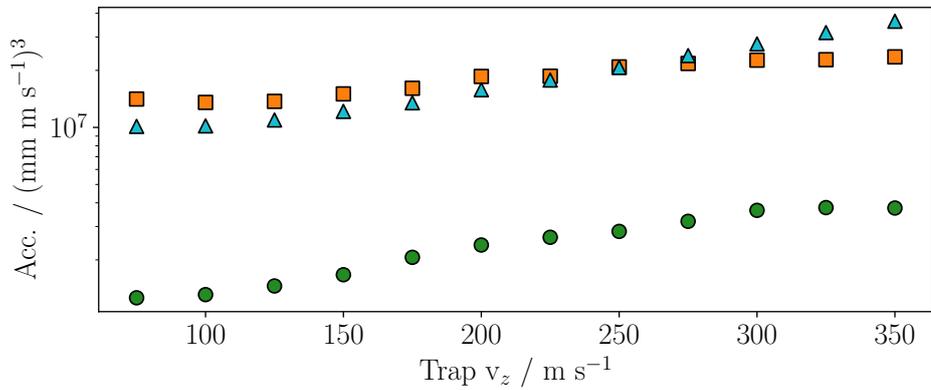


Figure 4.40: Estimated acceptance for hydrogen in deceleration mode for the 4-module machine. As before, the circles / squares / triangles are for PWM mode with no / wire / permanent magnet quadrupole. Compared to argon, the acceptance is much greater for this lighter species.

The code was adapted to model the  $m_F = 1$  sub-level of ground-state hydrogen. The results of deceleration of this state in the four-module decelerator are shown in figure 4.40; for comparison with argon these simulations all had an initial beam flow velocity of  $350 \text{ m s}^{-1}$ . It can be seen that for this much lighter species even the shorter decelerator could potentially remove the majority of the forward velocity of the beam. Regarding the two methods of estimating the acceptance, the ‘static’ method given in section 4.3 and the ‘dynamic’ one described here; the static method requires much less computation time, and can be run for an arbitrarily long simulated time to allow non-phase-stable particles to escape, and thus approaches the ideal or ‘true’ acceptance of the decelerator. However, it does not include effects arising from the PWM method, which the dynamic code does. The dynamic code, modelling the real physical experiment, does not allow sufficient time for all unstable atoms to escape, and so over-estimates the acceptance. Regardless, the results of both methods appear to be compatible to within an order of magnitude, and should therefore be useful in characterising the experiment. The discrepancy between the results for the alternative quadrupoles will be considered further in section 4.5.2.

## 4.5 Extending the Decelerator Simulations

Chapter 3 described the simulation codes we have developed in order to model the operation of the travelling-wave Zeeman decelerator, and gave examples of their application in reproducing experimental data. The codes can also be used to explore different combinations of parameters without the need to run the physical experiment, as well as to investigate possible modifications to the apparatus. Some of these results will be shown in the following sections, to begin we will extend our results for the deceleration of metastable argon.

### 4.5.1 Argon

Section 3.10 compared simulations with several sets of experimental data, for guiding and deceleration of metastable argon. Two data sets with a decelerator initial target trap velocity of  $330 \text{ m s}^{-1}$  were considered; in both cases the time delay between the source trigger and the activation of the power electronics was set to  $796 \mu\text{s}$ . This value was chosen according to an earlier estimate of the time taken for a synchronous atom moving at the trap velocity to reach the centre of the trap as it formed, and when we applied the methods described in section 4.1.1 they suggested that  $829 \mu\text{s}$  would be more appropriate. Figure 4.41 shows the result of a simulation with this delay but all other parameters the same as those given in section 3.10; *i.e.* deceleration from  $329 \text{ m s}^{-1}$  to  $321 \text{ m s}^{-1}$  with a peak current of  $400 \text{ A}$ . The target trap appears to be well loaded compared to that in figure 3.33; the shape of the loaded region of the phase-space plot suggests a more uniformly-filled longitudinal separatrix.

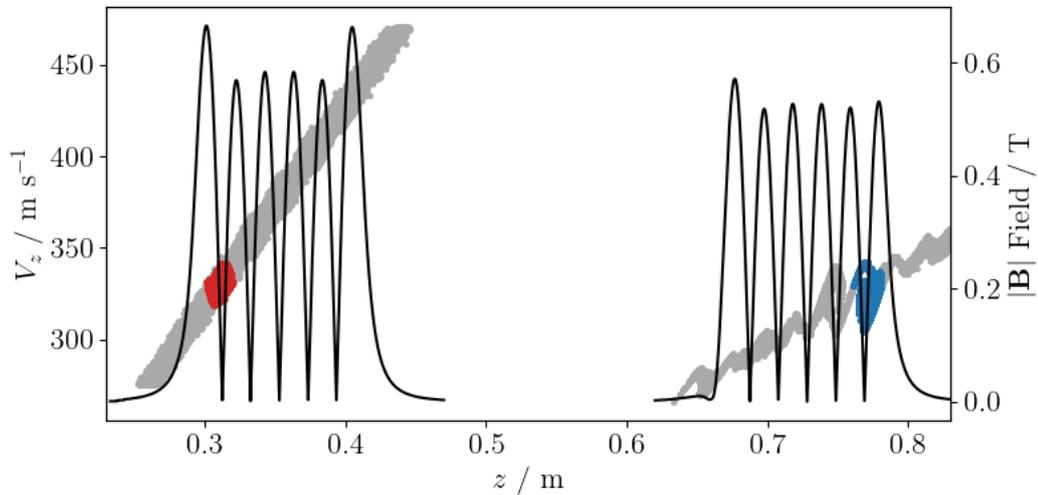


Figure 4.41: Longitudinal phase-space plots of attempted argon deceleration from  $329 \text{ m s}^{-1}$  to  $321 \text{ m s}^{-1}$ , peak current of  $400 \text{ A}$ , equivalent to figure 3.33 but with an electronics delay of  $829 \mu\text{s}$  rather than  $796 \mu\text{s}$ . The atoms located within the trap at the end are highlighted in blue, and their positions at the start in red.

The velocities of the spin-stretched  $m_J = 2$  atoms that reached the detector before (red) and after (blue) the deceleration are shown in figure 4.42; left is the whole sample, right is only the atoms located in trap 2. A deceleration of atoms is visible, with the mean velocity of the atoms in trap 2 decreasing from  $331$  to  $323 \text{ m s}^{-1}$ ; however, the allegedly less optimal equivalent in section 3.10 figure 3.34 had initial / final mean velocities of  $336 / 325 \text{ m s}^{-1}$ . An explanation can be found

by comparing figures 3.33 and 4.41; the less well-loaded trap in the earlier example effectively contained two separate populations of atoms with different central velocities, and not enough time was available during the deceleration process for them to merge, with the non-phase-stable atoms escaping the trap. Whilst some atoms had been decelerated, the values placed on the initial and final velocities were somewhat misleading. The example shown here has a well-populated phase-stable region so that the atoms have a rather more unimodal distribution, which maintains its approximate shape as the mean value is reduced during deceleration. More significantly, with the corrected delay setting the ratio of  $m_J = 2$  atoms in trap 2 compared to trap 3 at the end of the sequence was  $\approx 2:1$ , compared to  $\approx 1:4$  for the badly loaded simulation in section 3.10.

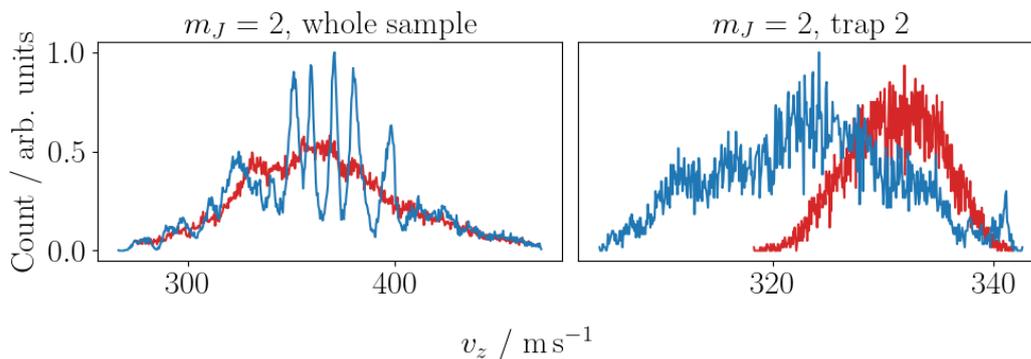


Figure 4.42: Comparison of initial (red) to final (blue) velocity distributions for the  $m_J = 2$  argon atoms that survive to reach the detector. Left, all atoms in the sample, right, only those located in the second trap at the end of the deceleration sequence. Comparable to figures 3.31 and 3.34 but with a more optimal power electronics delay for loading.

Again considering the badly-loaded example from section 3.10, we can explore the effects of adding an additional transverse quadrupole field, which was not present in the experimental data. Adding the wire quadrupole running at a current of 700 A we obtain phase-space and velocity distributions practically indistinguishable from figures 3.33, 3.34, 3.35 and 3.36, with the mean initial and final velocities of the atoms located in traps 2 and 3 differing only by  $\approx 1\%$ . This is unsurprising since the quadrupole guide has no components in the beam direction  $z$ . However, the number of  $m_J = 2$  atoms present in both trap 2 and 3 is increased by  $\approx 50\%$ .

We are also able to simulate increasing the number of deceleration coil modules from the existing 4 to the proposed 9 (see section 5). Compared to the real data, simulations of a 9-module machine set to decelerate from 330 to 320  $\text{m s}^{-1}$  (again with no quadrupole) show equivalent redistribution of the atomic velocities to within  $\approx 1\%$ , but with a reduction of those reaching the detector of  $\approx 25\%$ . The gentler deceleration achievable with a longer decelerator increases the acceptance of the phase-stable region but this is balanced by the increased time the particles have to drift laterally to the sample delivery tube and be lost, indicating the importance of the additional quadrupole guide.

Combining these ideas, we can attempt to determine an optimal deceleration of argon. The simulation shown in figure 4.43 utilised the 9 module decelerator operating at a peak current of 500 A, driven by 650 V. The wire quadrupole was included, operating at 700 A, and the power electronics delay was set at 822  $\mu\text{s}$ . The atom source characteristics were once again based on the experimental data described in section 3.10. The initial travelling wave velocity was set to 330  $\text{m s}^{-1}$  for comparison with these deceleration data, but the final velocity was chosen to be 290  $\text{m s}^{-1}$ . For these parameters the PWM achieved initial / final wave velocities of 330 / 292  $\text{m s}^{-1}$ .

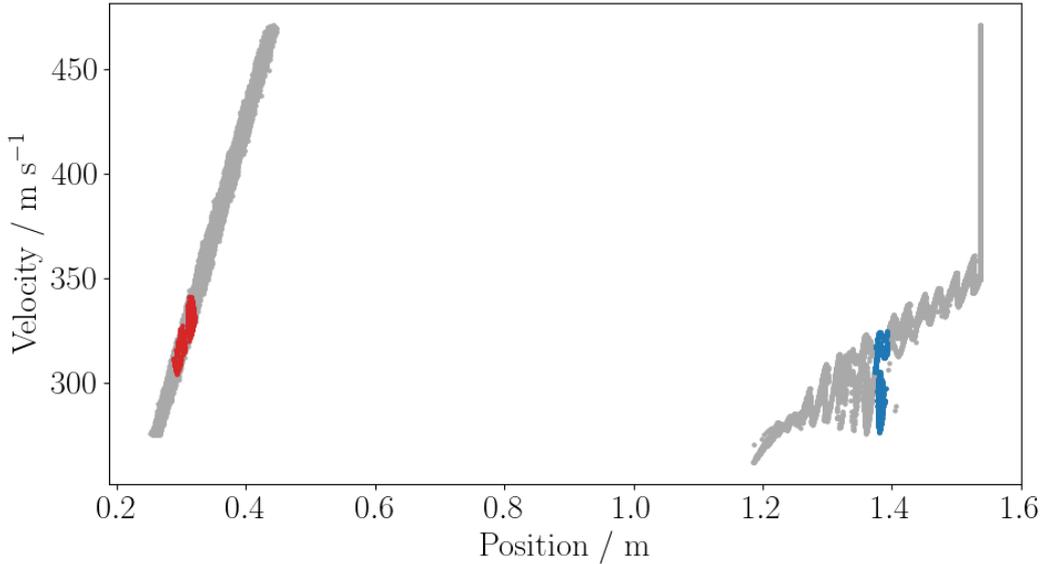


Figure 4.43: Longitudinal phase-space plots of attempted argon deceleration from 330 to 292  $\text{m s}^{-1}$ , peak current of 500 A, wire quadrupole active, proposed 9 module decelerator. Atoms shown are in the  $m_J = 2$  state.

Much clearer deceleration results are obtained for this configuration, with the mean longitudinal velocity of the spin-stretched  $m_J = 2$  atoms present in trap 2 as the field shuts down decreasing from 330 to 295  $\text{m s}^{-1}$ . These values are, however, somewhat distorted by the presence of a population of atoms that are not phase-stable and are leaving trap 2, identifiable in figures 4.43 and 4.44.

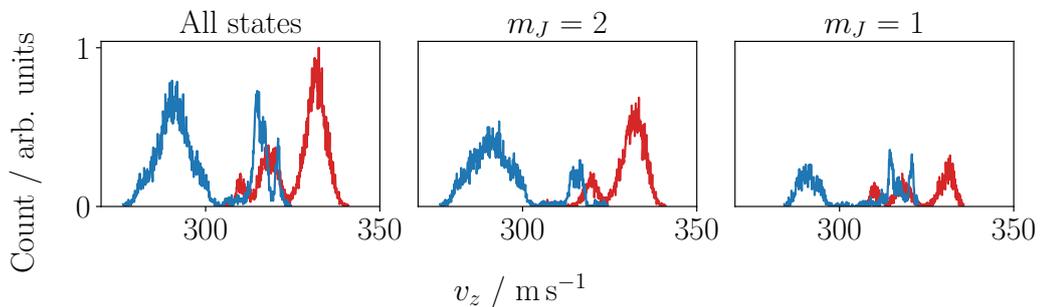


Figure 4.44: Velocities of the simulated argon atoms that remain in trap 2 at the end of the deceleration sequence from 330 to 292  $\text{m s}^{-1}$ . Red / blue are initial / final velocities. Both low-field-seeking states are decelerated, although there are twice as many  $m_J = 2$  than  $m_J = 1$  atoms.

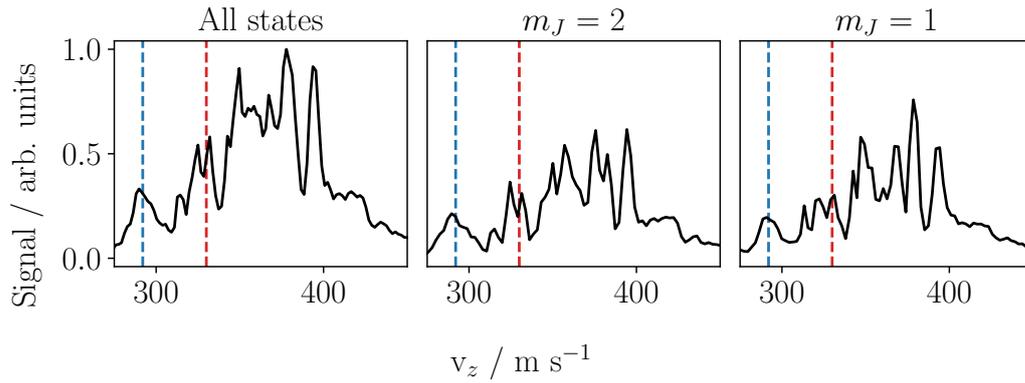


Figure 4.45: Deceleration of argon from  $330$  to  $292$   $\text{m s}^{-1}$ , proposed 9 module decelerator with quadrupole,  $500$  A at  $650$  V. Shown are the velocities of all magnetic states and the two low-field-seeking states. Red / blue lines indicate the initial and final velocities of the decelerator wave. See figure 3.30 for an equivalent view of real data.

As well as attempting to capture atoms with a mean velocity well below the central velocity of the beam, section 3.10 gave an example in which the decelerator wave was set to initially match it. For the data shown in figures 3.41 to 3.43 the decelerator field was set with initial / final velocities of  $360 / 350$   $\text{m s}^{-1}$ . Considering the proposed 9-module extended experiment, we can assess the extent to which it could successfully decelerate argon. Figures 4.46 and 4.47 show a simulation of deceleration from  $360$  to  $330$   $\text{m s}^{-1}$ , peak decelerator current is  $500$  A and the wire quadrupole is active. The initial atom distribution is chosen to match the real transmission data used in section 3.10. Much more significant spatial bunching of the atoms is seen than was the case for the shorter decelerator. The initial / final longitudinal velocities of the  $m_J = 2$  atoms that remain bunched in trap 2 at the end of the sequence are  $356.9 / 334.7$   $\text{m s}^{-1}$ , and represent  $\approx 1.1$  % of the initial sample.

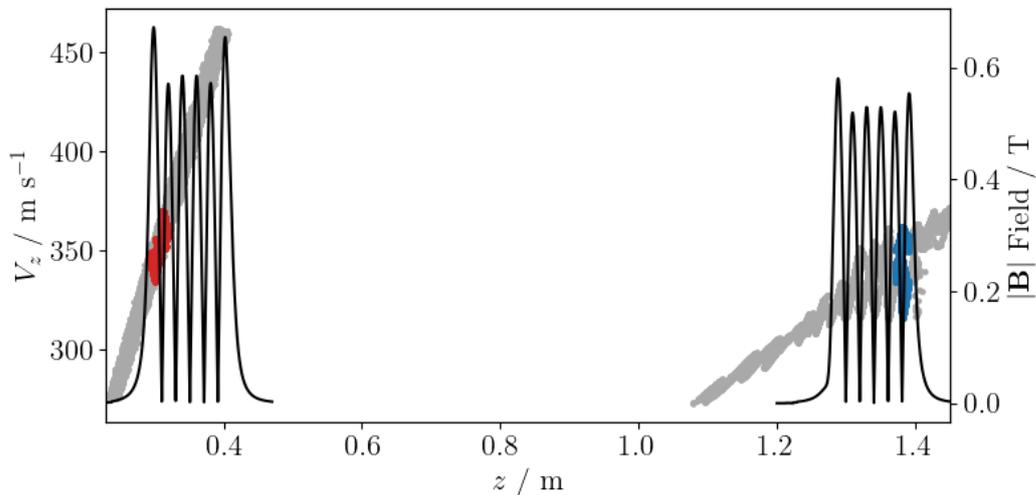


Figure 4.46: Longitudinal phase-space plots of attempted argon deceleration from  $360$   $\text{m s}^{-1}$  to  $330$   $\text{m s}^{-1}$ , peak current of  $500$  A, proposed 9-module decelerator. The atoms located within the trap at the end are highlighted in blue, and their positions at the start in red.

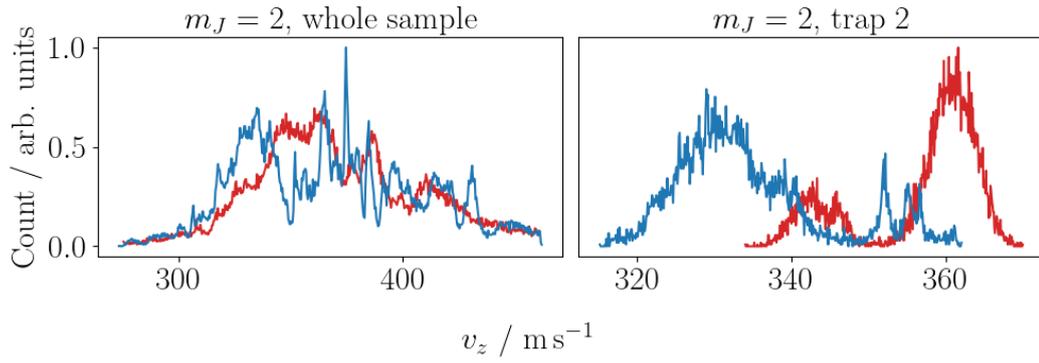


Figure 4.47: Argon deceleration from  $360 \text{ m s}^{-1}$  to  $330 \text{ m s}^{-1}$ , peak current of  $500 \text{ A}$ , proposed 9-module decelerator. Comparison of initial (red) to final (blue) velocity distributions for the  $m_J = 2$  atoms that survive to reach the detector. Left, all atoms in the sample, right, only those located in the second trap at the end of the deceleration sequence.

A harder deceleration still is shown in figures 4.48 and 4.49, with initial / final wave velocities of  $360 / 300 \text{ m s}^{-1}$ . The other parameters are unchanged from the previous case. The initial / final mean longitudinal velocities of the  $m_J = 2$  atoms that remain bunched in trap 2 at the end of the sequence are  $343.1 / 324.6 \text{ m s}^{-1}$ , and represent  $\approx 0.8 \%$  of the initial sample. The optimally-loaded atoms are seen to be decelerated but the distribution is not as clear as that of the less-harsh deceleration in the previous example. The mean velocities are clearly somewhat misleading here; there are multiple peaks distorting the average. However, there are strong peaks for the initial and target velocities.

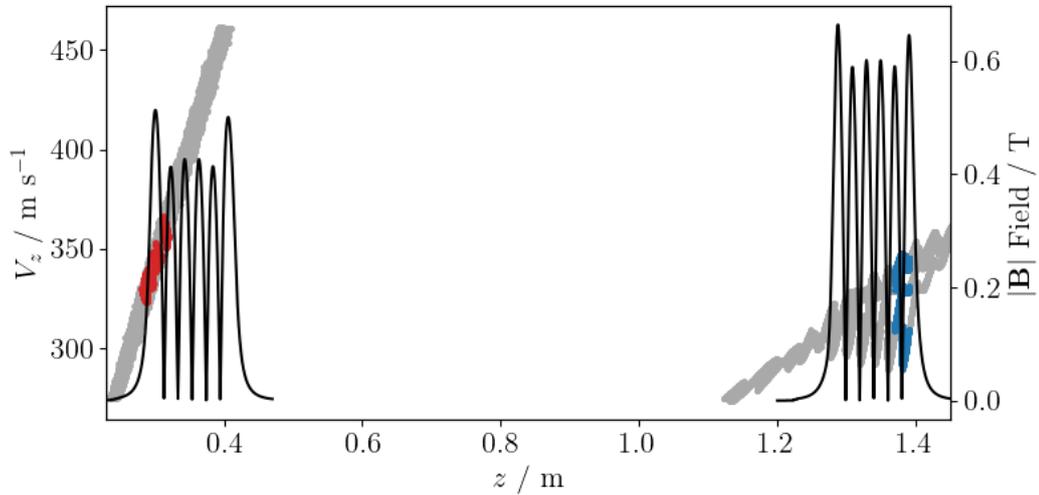


Figure 4.48: Longitudinal phase-space plots of attempted argon deceleration from  $360 \text{ m s}^{-1}$  to  $300 \text{ m s}^{-1}$ , peak current of  $500 \text{ A}$ , proposed 9-module decelerator. The atoms located within the trap at the end are highlighted in blue, and their positions at the start in red.

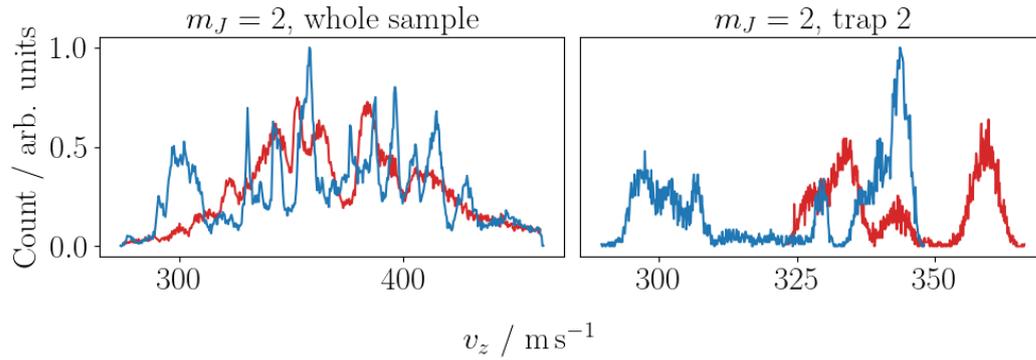


Figure 4.49: Argon deceleration from  $360 \text{ m s}^{-1}$  to  $300 \text{ m s}^{-1}$ , peak current of  $500 \text{ A}$ , proposed 9-module decelerator. Comparison of initial (red) to final (blue) velocity distributions for the  $m_J = 2$  atoms that survive to reach the detector. Left, all atoms in the sample, right, only those located in the second trap at the end of the deceleration sequence.

The multiple peaks present in the velocity distributions of the atoms located within the bounds of a single trap that are seen in *e.g.* figures 3.40 and 4.49 can be understood by taking a close look at a plot of longitudinal velocity versus position for the atoms towards the end of the deceleration sequence. This is shown in figure 4.50; the on-axis decelerator field magnitude is included to show the positions of the traps. The decelerator electronics delay setting was chosen to target trap 2, which is the trap located to the right at  $z \approx 1.38 \text{ m}$ , the one to the left at  $z \approx 1.36 \text{ m}$  is trap 3, and so on. It can be seen that whilst traps 2 and 3 have bunched atoms suggesting reasonably well-loaded separatrices, with central velocities in line with the trap velocity, atoms that are moving too fast to be trapped are skipping between traps. When analysing the effects of the decelerator field on the atoms by tracking the ones located in a particular moving trap, the presence of these atoms can distort the conclusion drawn. These effects can be minimised in several ways, where feasible; setting a gentler deceleration, using a colder source, using a lighter species, or increasing the trap depth by choosing a larger peak current.

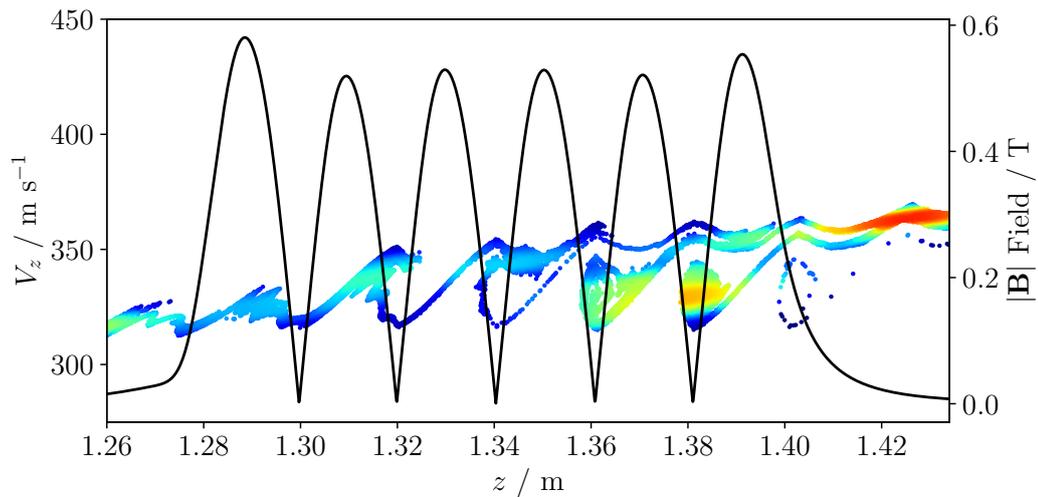


Figure 4.50: Argon deceleration  $360$  to  $330 \text{ m s}^{-1}$ , peak decelerator current is  $500 \text{ A}$ . The longitudinal positions and velocities of the atoms is presented with the axial decelerator field magnitude at the end of the deceleration sequence, as the field is just about to shut down. The shading gives the relative density of atoms, with red being the most dense.

To summarise this section; we have demonstrated some predictions for the results we would expect to get when decelerating metastable argon if various experimental parameters were adjusted, or if the physical apparatus was upgraded. Although argon is not of particular scientific interest to us, it has been chosen here for comparison to the characterisation data that we presented in chapter 3. Predictions regarding other species will be made in section 4.6 and chapter 5. In this section fairly conservative values of the peak decelerator current have been chosen, again for comparison to existing data, but higher currents are also explored in chapter 5. Our simulations suggest that for the existing 4-module decelerator and source, with an ideal set of experimental parameters we should be able to effectively decelerate a portion of a beam of argon from  $\approx 350$  to  $\approx 325 \text{ m s}^{-1}$  or lower, but this remains to be tested.

### 4.5.2 Assessing Quadrupole Effects

The decelerator fields are described in section 2.3.4, and it was shown that they are not strongly confining in the directions orthogonal to the beam axis, particularly along the  $x$ -axis. Additional transverse focussing is desirable to mitigate particle losses in these directions, and section 2.3.7 described two possible schemes. The experiment was initially constructed with a wire quadrupole comprising a high-current DC circuit, typically operating at  $\approx 700$  A (for specific details see McARD [2]). This design has undesirable fringe fields arising from the wire connectors at the ends, necessitating precise and complex control electronics in order to ensure they do not impinge upon the beam. Otherwise, it produces a strong, uniform quadrupole field in the  $xy$ -plane, but as shown in figure 2.15 it is not ideal when combined with the decelerator fields, so that the traps in the travelling wave are alternately weakened or strengthened. Extending the wire quadrupole to a longer decelerator would also be problematic, and so when a part of the system failed it was decided not to repair it, but to instead replace it with a permanent magnet guide. This design was expected to be superior to the wire guide for reasons additional to the lack of electronics; the magnetic field gradients are rather steeper, giving stronger confinement, and the problem of weakened alternating traps is avoided due to the field lines entering the trap along different axes. These differences are illustrated in figures 2.16 and 2.17.

Section 4.3 described a method of estimating the acceptance of a trap in the decelerator wave by modelling the motion of an initially-uniform sample of atoms in the trap. A pseudoforce can be added to represent deceleration effects. This was done for metastable argon with an initial beam velocity of  $350 \text{ m s}^{-1}$ , for both the 4-module and proposed 9-module decelerator, and the results are shown in figure 4.31. As expected, the presence of either quadrupole gives higher acceptance than the decelerator fields alone, with the wire quadrupole performing slightly worse than the permanent one. The fully dynamic acceptance model described in section 4.4 was used to repeat these tests. Figures 4.38 and 4.39 show the 4- and 9-module machine in deceleration mode, with results similar to those of the static simulation for higher final velocities. Somewhat surprisingly, however, for harder deceleration the performance of the permanent quadrupole dips below the wire.

The obvious conclusion when looking at figures 4.38 and 4.39 is that the estimated acceptance for the different quadrupoles is related to the time that the particles spend interacting with the trapping fields; namely, that the permanent quadrupole provides superior confinement for a short time, but the wire system is better for slower beams / longer decelerators. However, the static acceptance model is run for an artificially long time, allowing a closer approximation of the ‘true’ acceptance to be determined, albeit without including the effects of the oscillations in the trap field (see section 4.2). Therefore, the differing behaviour of the quadrupoles when compared to the static case is not simply due to the time that the atoms spend in the field, but is due rather to the nature of the decelerator fields in the fully time-dependent model.

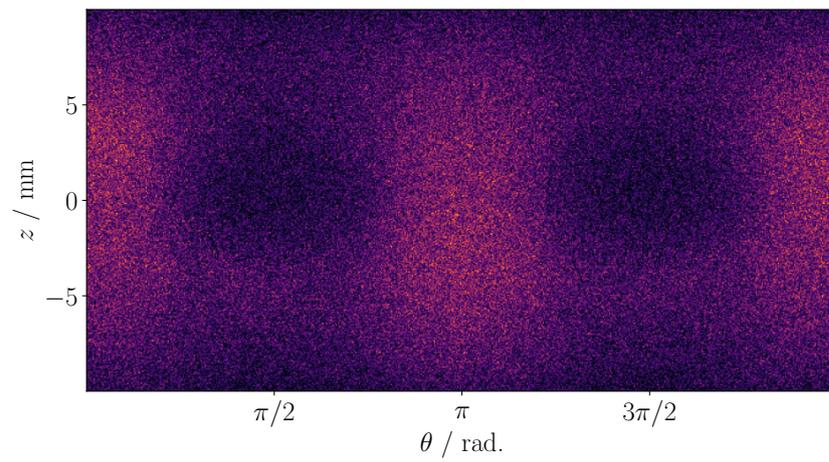


Figure 4.51: Argon loss in guiding mode, static simulation with decelerator coils at 500 A peak and wire quadrupole active at 700 A. A single trap in the decelerator field is modelled with an initially uniform sample of argon atoms as described in section 4.3. The density plot indicates atom losses at the inner surface of the sample delivery tube, depicted in cylindrical polar coordinates. The vertical axis is bounded by the trap longitudinal maxima, with the centre of the trap at  $z = 0$ . The angle  $\theta$  is defined such that  $\theta = 0$  and  $\theta = \pi$  lie along the  $x$ -axis.

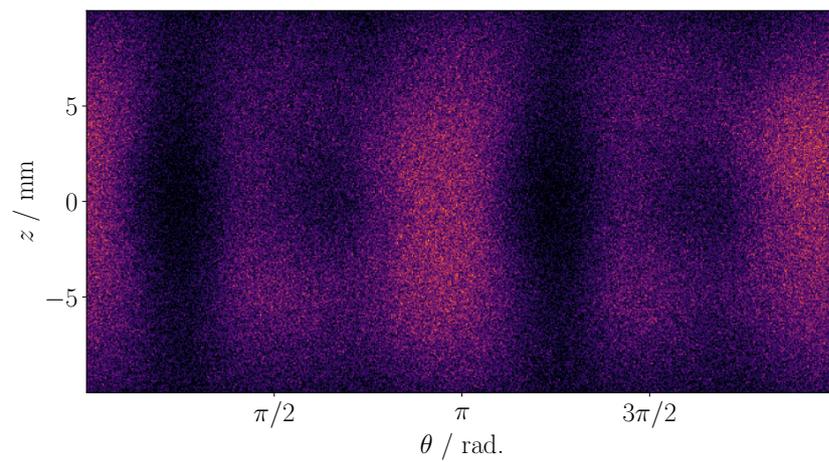


Figure 4.52: As per figure 4.51 but with permanent magnet quadrupole.

The static acceptance simulations were adjusted to record the leaving direction of escaping atoms, in order to verify that there was not a fundamental weakness in the design of the permanent magnet quadrupole. The code marks particles that have drifted to a radius  $r$  corresponding to the inner surface of the decelerator sample delivery tube as lost; figures 4.51 and 4.52 show density plots of these losses. The diagrams are converted into cylindrical polar coordinates, with the  $z$ -axis as the beam direction, centred on the trap and bounded by the forward and rear trap maxima.  $\theta = 0$  and  $\theta = \pi$  represent the points  $x = r, y = 0$  and  $x = -r, y = 0$ , respectively.

The equivalent plot for the system with no quadrupole is not shown, because it has little to no structure. The wire system, however, does show some structure, with the main regions of loss being in the  $x$  and  $-x$  directions. There are clear regions with little loss in the  $y$  and  $-y$  directions in the centre of the trap, reflecting the strong  $y$  focussing. It appears that atoms either leave along  $x$  or, if they move forward or back from the centre, can be guided to escape in  $y$ . The rear trap wall is slightly deeper than the front (see figure 4.33) so that the greater losses to the rear are somewhat surprising. Overall, the potential has a complicated shape, and taking views along given axes or planes does not allow us to easily predict the behaviour of all plausible particle trajectories. This illustrates the importance of such Monte Carlo simulations.

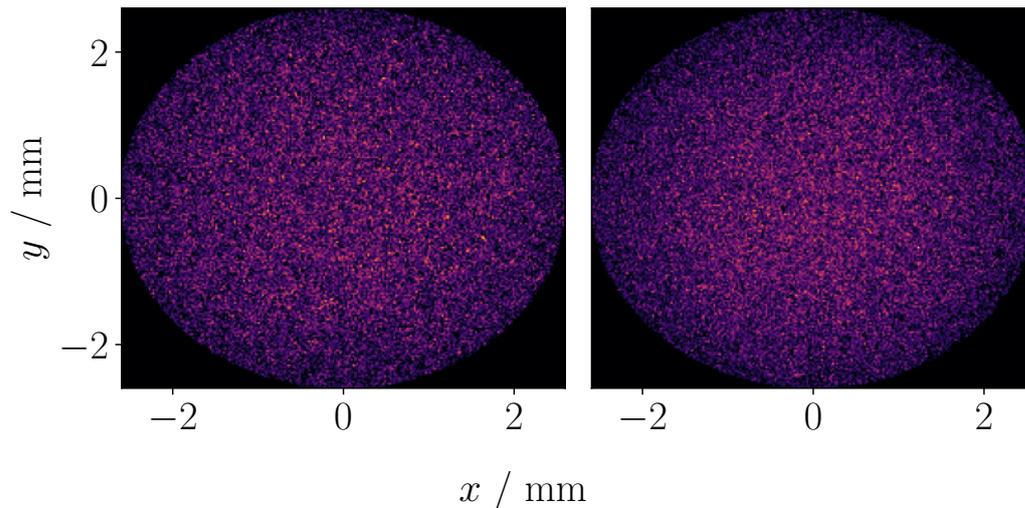


Figure 4.53: Argon loss in guiding mode, static simulation with decelerator coils at 500 A peak. Left, wire quadrupole active at 700 A, right, permanent magnet quadrupole. Shown are the  $xy$  positions of atoms at the instant they escape by passing the forward maximum at  $z \approx 10$  mm.

The permanent magnet quadrupole model shows behaviour that at first sight appears rather different from the wire system. Closer inspection shows that the losses in the  $x$ -directions have a similar structure. There are reduced losses along one diagonal plane, although more atoms have escaped in  $y$  than for the wire case. The typical asymmetry of the decelerator field plus permanent quadrupole along the  $x = y$  and  $x = -y$  axes, as shown in figure 2.16, can be clearly seen. The structure of this plot is suggestive of the pattern shown in figure 4.51 if it was ‘squeezed’ along one diagonal.

The static field simulation used here was run for a simulated time of 3 ms, which is rather longer than the atoms would spend in the decelerator field for the the current 4-module machine at a typical wave velocity, but this is necessary to allow time for the non-phase-stable atoms to escape. For this example of constant-velocity guiding the permanent quadrupole resulted in  $\approx 9\%$  more retained atoms than the wire, although both systems had over a factor of five more atoms at the end than the simulation with no quadrupole. More revealingly, the permanent quadrupole actually lost  $\approx 10\%$  fewer atoms in the  $xy$  direction than the wire, but more in  $z$ . As seen in figure 4.53, plotting the pattern of final coordinates for atoms reaching the forward trap maximum at  $z \approx 10$  mm showed no significant structure for either system, apart from a hint that the losses from the trap with the permanent quadrupole were more strongly clustered at smaller radii from the beam axis. To further investigate, a repeat of these simulations under significant deceleration was performed.

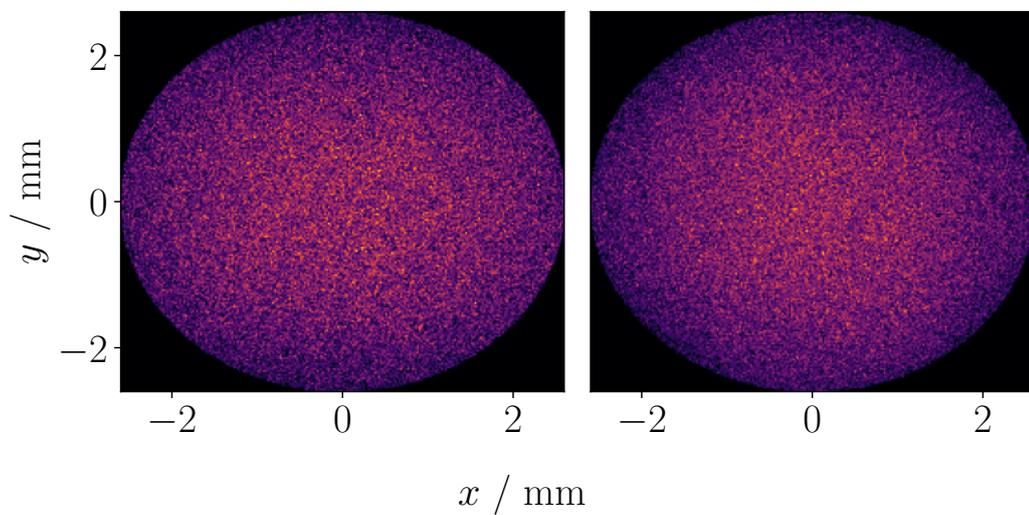


Figure 4.54: Argon deceleration from  $350$  to  $275 \text{ m s}^{-1}$ , static simulation with decelerator coils at  $500 \text{ A}$  peak. Left, wire quadrupole active at  $700 \text{ A}$ , right, permanent magnet quadrupole. Shown are the  $xy$  positions of atoms at the instant they escape by passing the forward maximum at  $z \approx 10$  mm.

Figure 4.54 shows the atoms escaping in the  $z$ -direction for deceleration from  $350$  to  $275 \text{ m s}^{-1}$ . As expected, both plots show higher intensity of atoms than in the guiding case, since the deceleration pseudoforce is driving them forwards. Both plots show that the atoms are more dense closer to the beam axis at the origin, with the permanent quadrupole displaying a suggestion of the asymmetry along  $x = y$  and  $x = -y$  that was alluded to earlier. Neither show a significant weak point or ‘hole’ through which atoms are being lost.

Figures 4.55 and 4.56 give the losses in the  $xy$ -direction. Compared to figures 4.51 and 4.52 both systems show, as expected, that the losses are now more significant for positive  $z$  as the atoms are pushed to the front. The overall patterns of loss are otherwise very similar to those shown for the guiding case. For this harsh test, modelling hard deceleration in the 4-module decelerator, all atoms were lost from the trap region after 3 ms. Comparing the directions of loss, the wire system lost  $54\%$  in  $xy$  and the remainder in the positive  $z$ -direction, compared to  $48\%$  for the permanent quadrupole.

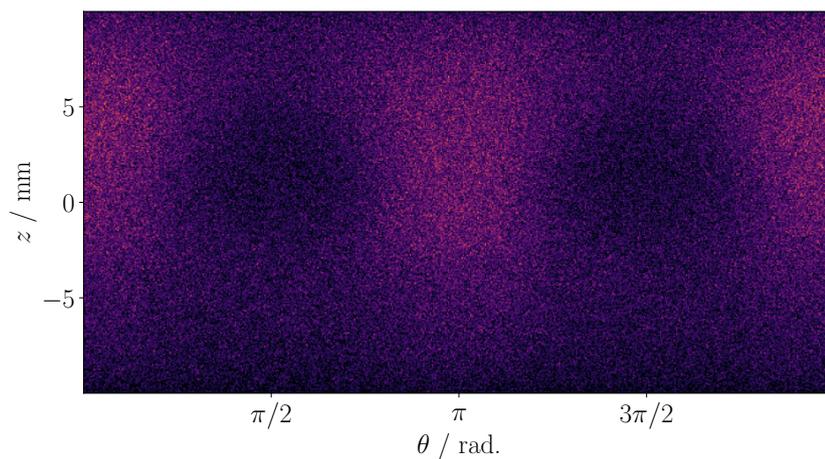


Figure 4.55: Argon deceleration from  $350$  to  $275 \text{ m s}^{-1}$ , static simulation with decelerator coils at  $500 \text{ A}$  peak and wire quadrupole active at  $700 \text{ A}$ . A single trap in the decelerator field is modelled with an initially uniform sample of argon atoms as described in section 4.3. The density plot indicates atom losses at the inner surface of the sample delivery tube, depicted in cylindrical polar coordinates. The vertical axis is bounded by the trap longitudinal maxima, with the centre of the trap at  $z = 0$ . The angle  $\theta$  is defined such that  $\theta = 0$  and  $\theta = \pi$  lie along the  $x$ -axis.

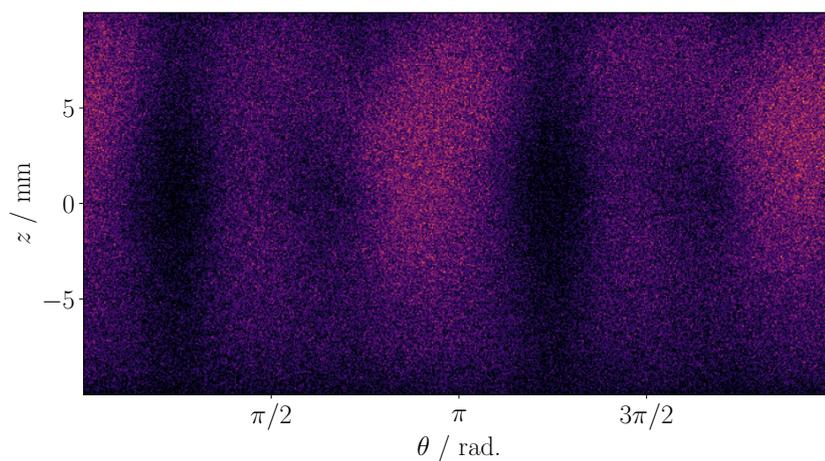


Figure 4.56: As per figure 4.55 but with permanent magnet quadrupole.

The code to count the atom leaving directions was then added to the dynamic acceptance code (see section 4.4). Figure 4.57 shows the percentage particle loss in the radial, forward and rear directions of a single trap in the decelerator field as a function of guiding velocity.

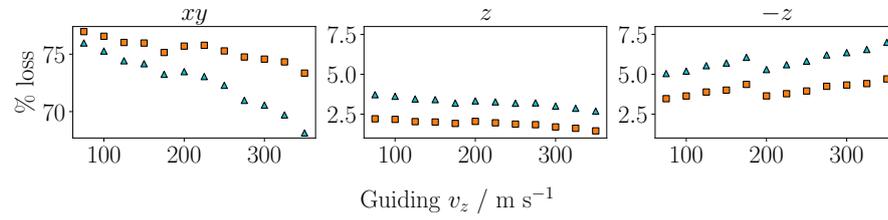


Figure 4.57: Percentage loss in the (left) radial, (centre) forward and (right) rear longitudinal directions, time-dependent acceptance simulations. Constant-velocity guiding mode, decelerator fields set to peak 500 A. Square, wire quadrupole, triangle, permanent magnet.

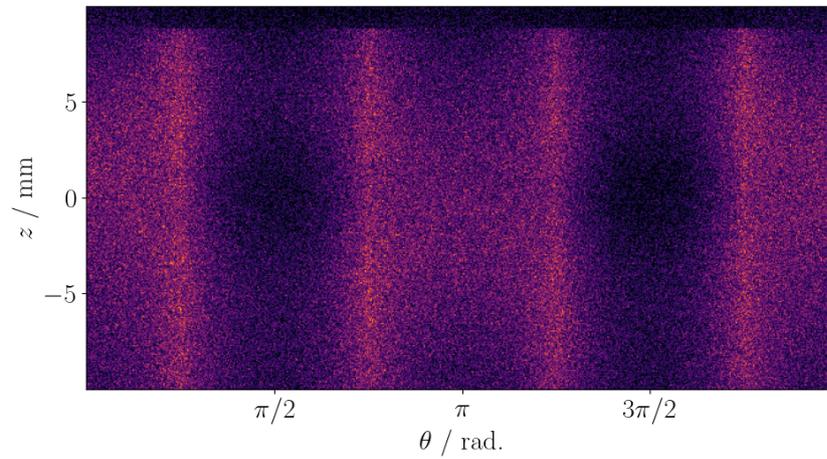


Figure 4.58: Argon loss in guiding mode at  $150 \text{ m s}^{-1}$ , dynamic simulation as described in section 4.4, decelerator coils at 500 A peak and wire quadrupole active at 700 A.

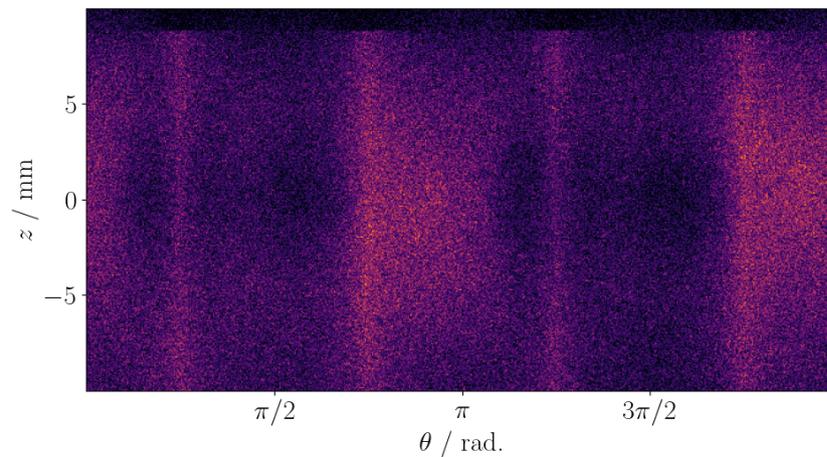


Figure 4.59: As per figure 4.58 but with permanent magnet quadrupole.

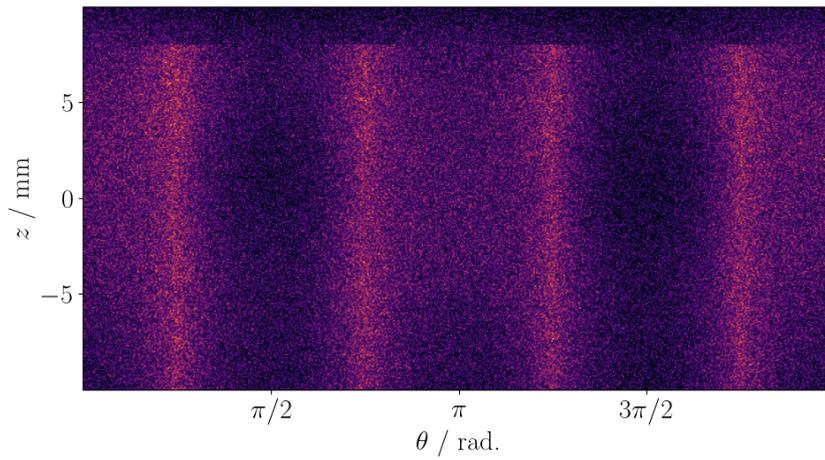


Figure 4.60: Argon loss in deceleration mode from  $350$  to  $275 \text{ m s}^{-1}$ , dynamic simulation as described in section 4.4, 4-module decelerator coils at  $500 \text{ A}$  peak and wire quadrupole active at  $700 \text{ A}$ .

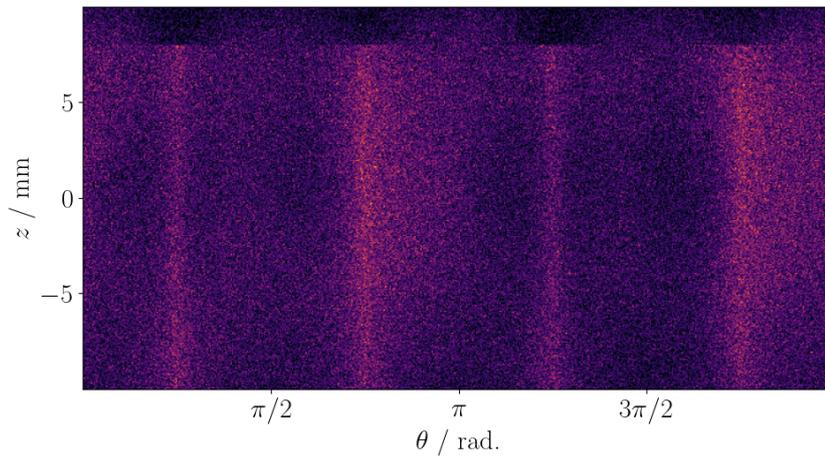


Figure 4.61: As per figure 4.60 but with permanent magnet quadrupole.

Figures 4.58 and 4.59 show the particle losses in the dynamic code, when guiding at  $150 \text{ m s}^{-1}$ . The time variation of the decelerator fields is driving the direction of the radial losses to be along the diagonals, when compared to the static code. Running in deceleration mode further enhances this loss, as shown in figures 4.60 and 4.61. Figure 4.57 shows that the permanent quadrupole has significantly less loss radially, especially at higher wave velocities, but worse longitudinal losses. The different structure seen in the dynamic guiding and deceleration plots as compared to the static ones is probably not that significant, since it seems that the main driver of the difference in acceptance is losses in  $z$ . Further work is needed to investigate the loss channels in the front and rear of the decelerator traps. To some extent it may be the case that particles that are moving too fast to be trapped and would escape radially are, in the case of the permanent magnet system, deflected to escape longitudinally, but this cannot be the entire picture since the total losses are larger when compared to the wire quadrupole. Therefore, some combination of the permanent magnet and decelerator fields is driving particle loss in  $z$  even whilst losses in the  $xy$ -plane are reduced.

## 4.6 Comparison to Other Travelling-Wave Zeeman Decelerators

Section 2.3.3 gave a brief description of several other travelling-wave / moving-trap Zeeman decelerators, for convenience referred to as the ‘Paris’ [51], ‘Tel Aviv’ [52] and ‘Basel’ [53] experiments. In this section we attempt to compare some of the results of these experiments with the achieved and predicted results of the Durham decelerator.

### 4.6.1 Paris Decelerator

The most straightforward comparison is with the Paris decelerator, since the Durham machine was based upon it. Trimeche *et al.* demonstrated constant-velocity guiding of metastable argon at 430 and 464  $\text{m s}^{-1}$  [51]. The time-of-flight traces presented are very clear with significant peaks representing the atoms that are well confined in the moving trap of the decelerator wave. The data are also well reproduced by simulation, which allows insight into the positions and velocities of the particles in the beam. Although we also use argon we have no directly comparable data since the flow velocity of the Paris General valve source is estimated to be 465  $\text{m s}^{-1}$  with a translational temperature of 3 K when cooled to 100 K with a backing pressure of 2 bar. This is considerably faster than our Even-Lavie valve, which produces an argon beam at an estimated 370  $\text{m s}^{-1}$  / 7 K when cooled to 140 K at 5 bar. In order to simulate a faster beam source we generated a velocity distribution using equation 3.2.2, with  $v_T = 465 \text{ m s}^{-1}$  and  $\sigma_{T,z} = 25 \text{ m s}^{-1}$ . The spatial and temporal distribution of the atoms were chosen according to the geometry and typical discharge parameters of our source (as described in section 3.2). The simulation code was then configured for constant-velocity guiding at 465  $\text{m s}^{-1}$ , peak current of 350 A, wire quadrupole active at 700 A.

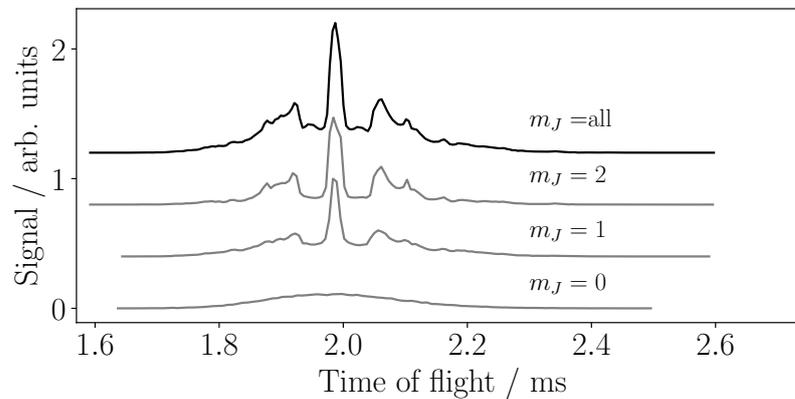


Figure 4.62: Simulated guiding of  $^3\text{P}_2$  argon, decelerator operating at 465  $\text{m s}^{-1}$  with peak current of 350 A. Black trace is all atoms, the labelled grey traces are the sub-levels with projection number  $m_J = 0, 1, 2$ . Parameters chosen to compare with figure 7 in Trimeche [51].

Figure 4.62 shows the simulated time-of-flight profile. Overall, the results have reasonably good agreement with figure 7 of Trimeche [51], having the same general shape and features. The greater length of our decelerator can be seen in the position of the peaks. The exact dimensions of the Paris experiment are not given, but it appears to be  $\approx 620$  mm from the valve nozzle to the detector, compared to 924 mm in Durham. The plot shows the overall trace and the contributing traces of the two low-field-seeking states  $m_J = 1, 2$  and the neutral state  $m_J = 0$ . Trimeche state that both the  $^3P_0$  and  $^3P_2$  metastable terms will be present with ‘all the Zeeman sub-levels equally populated’, but it is not clear what they mean by this - it could be interpreted to mean that  $1/5$  of the total population will be  $m_J = 0$ , or  $1/3$  of the total if the  $J = 0, m_J = 0$  and  $J = 2, m_J = 0$  populations are counted separately. If both  $^3P_0$  and  $^3P_2$  are assumed to be equally populated then  $3/5$  of the atoms will be in the  $m_J = 0$  state. The lower curve in the Trimeche figure is small compared to the others, suggesting that they have taken the first interpretation and included the five Zeeman sub-levels equally, and so our simulation has the same initial distribution. The Trimeche simulations are a good match to their recorded signal, but as discussed in chapter 3 we believe that our source is producing significantly more  $^3P_0$  than  $^3P_2$ , so that a simulated time-of-flight profile that does not include extra  $^3P_0$  does not well represent our data. To illustrate this, the plots in figure 4.63 show the same time-of-flight as figure 4.62 but with extra  $^3P_0$  added in a 5:1 ratio, and it can be seen that these results look more like those given in chapter 3.

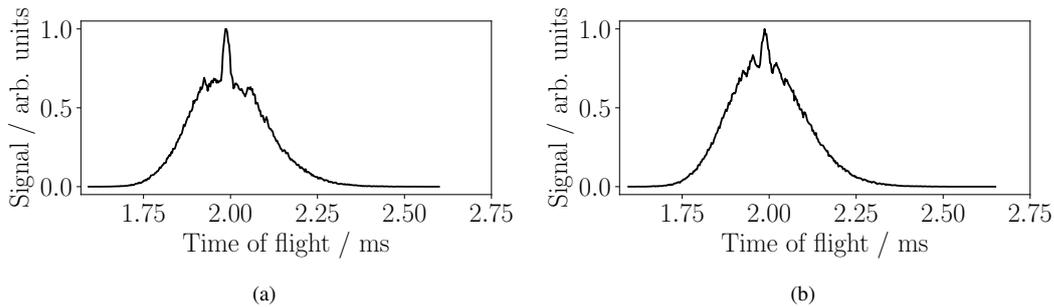


Figure 4.63: Simulated guiding of  $^3P_2$  argon as per figure 4.62, however extra atoms in the  $^3P_0$  term added in the ratio 5:1 with  $^3P_2$ . Left, wire quadrupole, right, no quadrupole.

Figure 4.64 shows a longitudinal phase-space plot of the  $m_J = 2$  atoms at the end of the guiding sequence; the atoms shown in blue are within the maxima of the moving trap, and the estimated separatrix (see section 4.1.2) is overlaid in black. This can also be compared to figure 7 in Trimeche; both the Durham and Paris decelerators show a similar spread of trapped velocities,  $\approx \pm 20$  m s $^{-1}$ , although somewhat larger for the Durham results due to the deeper traps. The different longitudinal dimensions of the Paris and Durham traps,  $\approx 6$  mm and  $\approx 20$  mm respectively, can be seen.

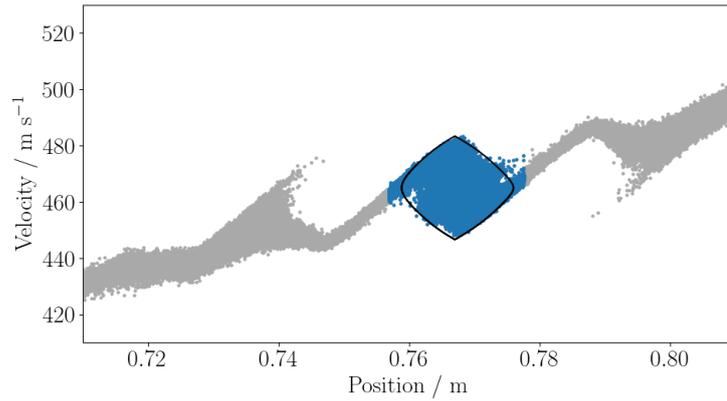


Figure 4.64: Simulated guiding of  ${}^3\text{P}_2$  argon as per figure 4.62. Longitudinal phase-space plot of atoms in the  $m_J = 2$  sub-level at the end of the deceleration sequence. Atoms within the trap maxima are in blue, the estimated separatrix is in black.

Sections 4.3 and 4.4 described methods for estimating the acceptance of the decelerator using either a static or time-dependent model of the trap fields. The former is not meaningful for guiding mode, since in the absence of deceleration pseudoforces the trapping potential is not velocity-dependent. The Trimeche results did not include an estimated acceptance, but we have calculated it for guiding at  $465 \text{ m s}^{-1}$  using the dynamic code model, for comparison with the results shown in figure 4.34 in section 4.4. We find that for no / wire / permanent magnet quadrupole we obtain  $7.9 \times 10^5 / 1.7 \times 10^6 / 3.0 \times 10^6 \text{ (mm m s}^{-1}\text{)}^3$ , compared to  $4.8 \times 10^5 / 1.2 \times 10^6 / 1.9 \times 10^6 \text{ (mm m s}^{-1}\text{)}^3$  when guiding at  $350 \text{ m s}^{-1}$ . These values can be somewhat misleading; the acceptance appears to increase with guiding speed because there is not enough time for phase-unstable particles to escape the trap.

Additionally, Bera reported guiding of argon at  $480$  and  $464 \text{ m s}^{-1}$  as well as deceleration from  $400$  to  $380$  and  $370 \text{ m s}^{-1}$ , and from  $395$  to  $365 \text{ m s}^{-1}$  [221], again seeing good agreement between data and simulations. Trimeche then reported deceleration of argon from  $464$  to  $420 \text{ m s}^{-1}$  and from  $430$  to  $402 \text{ m s}^{-1}$  [222]. Compared to the deceleration experiments reported in this thesis, which had a much lower beam flow velocity, these values are impressive. Repeating the caveat that we are using an approximate reproduction of the Paris initial beam conditions we have simulated deceleration of argon from  $464$  to  $420 \text{ m s}^{-1}$  with the Durham decelerator.

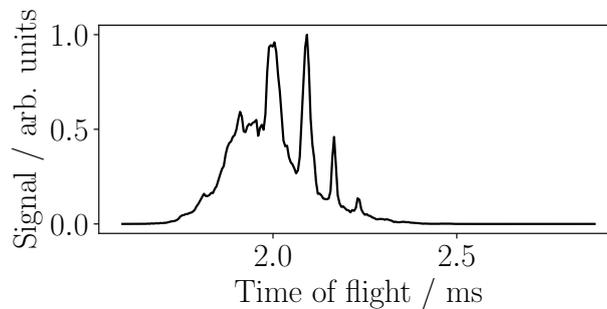


Figure 4.65: Simulated deceleration of  ${}^3\text{P}_2$  argon, decelerator operating from  $465$  to  $420 \text{ m s}^{-1}$  with peak current of  $350 \text{ A}$ , wire quadrupole active.

The time-of-flight in figure 4.65 does show a sharp peak at about the expected arrival time of atoms decelerated to  $\approx 420 \text{ m s}^{-1}$ , but the phase-space plot in figure 4.66 shows that there are few atoms in the central region of the trap. The shape of the atom sample in the centre of the trap suggests that it is not phase-stable; *i.e.* the atoms are only present in the trap because they have not yet had time to escape. To confirm this we attempted to plot the longitudinal separatrix using the method from section 4.1.2, but found no phase-stable region in the trap when such a large deceleration pseudoforce was applied. For the Durham decelerator the necessary pseudoforce equates to a deceleration of magnitude  $\approx 4.1 \times 10^5 \text{ m s}^{-2}$ ; for the shorter Paris design there is an even larger deceleration of  $\approx 7.2 \times 10^5 \text{ m s}^{-2}$ . This corresponds to a loss of kinetic energy of  $\approx 18\%$ ; for the deceleration data presented in chapter 3 there was a reduction of  $\approx 6\%$ . In later experimental work we have demonstrated a reduction of 10 - 15% of kinetic energy of argon (see appendix A), which is close to the maximum that can be achieved for this species with our decelerator in its current form. The predictions made in this chapter suggest that for the proposed 9 module extension a maximum of  $\approx 30\%$  of the kinetic energy of an argon beam could be removed. The next section will discuss the possibility of extending the decelerator further still.

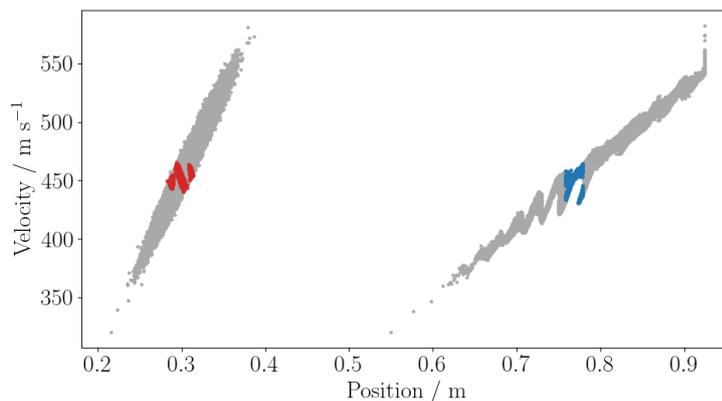


Figure 4.66: Simulated deceleration of  $^3\text{P}_2$  argon as per figure 4.65. Longitudinal phase-space plot of atoms in the  $m_J = 2$  sub-level at the end of the deceleration sequence, atoms located within the trap at the end are marked in blue, with their initial positions in red.

Despite their similarities, there are key differences in the Paris and Durham decelerators which must be borne in mind when comparing their experimental results. The most obvious difference seems to be related to the beam source; the Paris experiment appear to have been able to produce a beam of metastable argon that was largely composed of atoms in low-field seeking states by the time it reached the detector. This is most likely related to the different excitation schemes used. The tube through which the beam flows in the Paris experiment is 1.2 mm diameter, compared to 5.2 mm in Durham, so that particles that are not transversely focussed by the decelerator and quadrupole fields are more likely to be lost to collisions with the tube wall, but this is unlikely to be a very significant factor in the difference in Zeeman sub-level populations we see.

The distance from the end of the Durham decelerator to the detector is 140 mm, giving the guided or decelerated particles time to disperse along the propagation axis and thus broadening and weakening the peaks seen on the time-of-flight profile. This distance is even larger for the Paris experiment at 210 mm, but this is compensated for by the much faster beam flow velocity as well as by the more confined dimensions of the Paris moving traps. This can be seen by comparing the simulations in figure 4.63 with those in section 3.10. Also, the Paris excitation source produces pulses of  $\approx 2 \mu\text{s}$  duration, about an order of magnitude smaller than the temporal width of the samples we generate with our DBD (see section 2.2.3). Although this means that there are fewer excited atoms they have a much smaller temporal width, and it was therefore possible for the Paris experiment to effectively load only a single trap of their field. The faster beam and smaller dimensions also meant that the separatrix was more completely filled on loading than in our simulations.

Trimeche also describe the presence of secondary peaks which complicate the signal, caused by the interaction with the trapping fields of particles that are not phase-stable. They note that in a longer decelerator these secondary effects will be lessened as particles have more time to escape the unstable regions. Although our decelerator is longer we still have these issues due to our slower beam, which means that one or more traps adjacent to the target will be at least partially populated. We explored this in section 4.5, where we considered decelerating argon in an extended version of our machine. It will be further considered in chapter 5, when we simulate the deceleration of lighter species.

### 4.6.2 Tel Aviv Decelerator

The ‘Tel Aviv’ decelerator has a different design, as described in section 2.3.3, being composed of cylindrical coils arranged in overlapping anti-Helmholtz configuration. We will initially consider the results in Lavert-Ofir [52], in which deceleration of metastable neon was reported. This first version of the decelerator was 1.14 m in length, and although the coil geometry is very different from the Durham decelerator the source is similar; an Even-Lavie valve with dielectric barrier discharge. With the valve cooled to 74 K a beam velocity of  $430 \text{ m s}^{-1}$  and translational temperature of 399 mK was estimated. Deceleration data for a range of final velocities down to  $50 \text{ m s}^{-1}$  are presented, and reproduced in simulations which show excellent agreement with the measured time-of-flight profiles. Detection is via a microchannel-plate detector mounted on a translation stage, allowing measurement of the beam velocity and translational temperature as described in section 3.1.2. The acceptance is reported as being  $8 \times 10^6 / 4 \times 10^6 (\text{mm m s}^{-1})^3$  for final velocities of  $350 / 50 \text{ m s}^{-1}$ .

Although we did some exploratory work with metastable neon in the early phases of our experiment [1] we switched to argon in order to achieve slower beams. This is because in a supersonic expansion the flow rate has an inverse dependency on the particle mass (see equation 2.2.10 in section 2.2.1). We have used our static model code to estimate the acceptance when decelerating this species. The metastable  $^3\text{P}_2$  term of neon has the same Zeeman sub-level splitting as argon (shown in figure 3.14) so we can easily adapt our argon simulation code by adjusting the atomic mass. Figure 4.67 shows the results for an initial velocity of  $430 \text{ m s}^{-1}$ ; figure 4.68 is the same but for the proposed 9 module extended decelerator, since this would be of comparable length at 1.1 m. For both cases results were calculated with no quadrupole or with the wire or permanent magnet quadrupoles active.

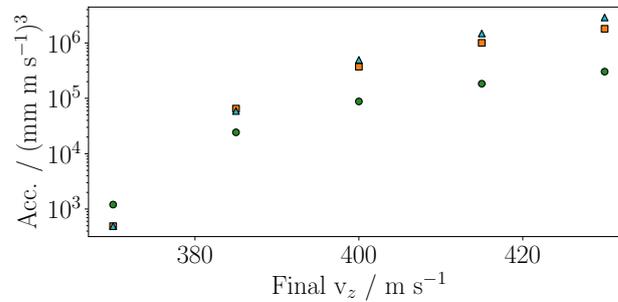


Figure 4.67: Estimated acceptance of the 4 module decelerator in deceleration mode, for the  $m_J = 2$  sub-level of  $^3\text{P}_2$  neon, peak current of 350 A. The circles / squares / triangles are for no / wire / permanent magnet quadrupole.

It can be seen that the Durham decelerator is not able to match the acceptance of the Tel Aviv decelerator. Sections 4.3 and 4.4 presented some estimates of acceptance for argon guiding (albeit with a lower initial velocity) and it is reasonable to expect that neon will have somewhat larger acceptance, since it has a more favourable magnetic moment to mass ratio. We see however that even considering the proposed extended decelerator with additional quadrupole, for deceleration from  $430$  to  $350 \text{ m s}^{-1}$  our acceptance is lower than the Tel Aviv results by an order of magnitude, and the performance drops rapidly for harder deceleration.

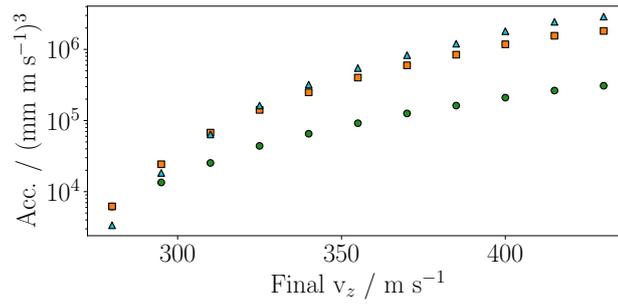


Figure 4.68: Estimated acceptance of the proposed 9 module decelerator in deceleration mode, for the  $m_J = 2$  sub-level of  ${}^3\text{P}_2$  neon, peak current of 350 A. The circles / squares / triangles are for no / wire / permanent magnet quadrupole.

Figure 4.69 shows phase-space plots of the initial and final atom distributions along the three axes, and can be compared to figures 2c and 2d in Lavert-Ofir [52]. The Durham spatial acceptance is about half that of the Tel Aviv machine in the radial directions, due to the smaller beam tube. The radial velocity acceptance is not symmetric about  $x$  and  $y$  in Durham due to the coil geometry, in  $y$  the Durham machine is only slightly weaker with  $\approx \pm 20 \text{ m s}^{-1}$ , although it only achieves half that in  $x$ . Longitudinally, the Tel Aviv traps are limited by the 10 mm distance between coil centres, which is about half the length of the Durham traps, but the Durham phase-stable region is smaller at about 8 mm and  $\pm 15 \text{ m s}^{-1}$  compared to 10 mm and  $\pm 25 \text{ m s}^{-1}$ . The Tel Aviv decelerator has approximately 25 % greater trap depth for this deceleration, at least in the forward direction; the asymmetry of the forward and rear walls of the trap can be seen in the shape of the separatrix.

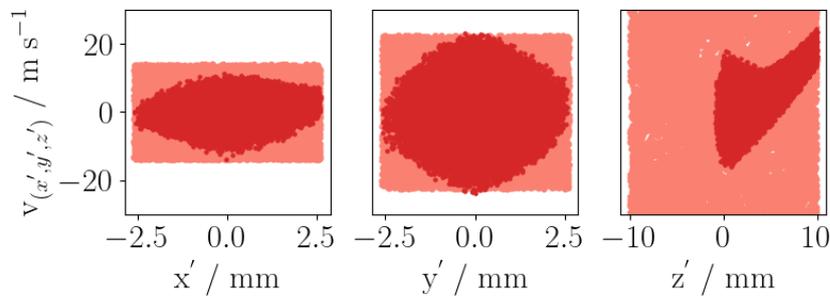


Figure 4.69: Estimation of decelerator phase-space acceptance for  ${}^3\text{P}_2$  neon by solving particle trajectories in a single trap of the decelerator, peak current of 350 A, wire quadrupole at 700 A, for deceleration from 430 to  $355 \text{ m s}^{-1}$ . An initial sample of atoms with a uniform spread of positions and velocities is created (light red) and allowed to evolve in the trap until only phase-stable atoms remain (dark red). The fraction that remains is multiplied by the initial phase-space volume to give an estimate of the acceptance.

As described in Akerman *et al.* [56], the Tel Aviv decelerator was then extended to 480 coil pairs spanning 2.4 m. They report deceleration of a dual-species beam comprising 1 % molecular oxygen and 99 %  ${}^3\text{P}_2$  metastable argon. The number of argon atoms in the spin-stretched  $m_J = 2$  state is increased by optical pumping at 811 nm, and a flow velocity of  $450 \text{ m s}^{-1}$  is reported. Results are presented for guiding at the beam mean velocity and deceleration to various velocities down to  $100 \text{ m s}^{-1}$ . The excited argon atoms can be directly detected by a microchannel-plate detector, the oxygen molecules are first ionized in a  $2 + 1$  resonance enhanced multi-photon ionization process.

Significant detected signals are visible even for deceleration to  $100 \text{ m s}^{-1}$ , which represents a reduction of  $\approx 95 \%$  of the kinetic energy of the particles.

We have a good understanding of the capabilities of the Durham decelerator when working with metastable argon, as discussed in chapters 3 and 4. However, even with the proposed 9 module extended decelerator we estimate that we could at best remove  $\approx 30 \%$  of the kinetic energy of an argon beam with an initial flow velocity of  $350 \text{ m s}^{-1}$ . For comparison with the data in Akerman we have used our static code model to estimate our acceptance for argon and oxygen if our decelerator was extended further, to 20 modules, which would provide a 2.4 m long deceleration stage.

Due to spin-rotation interaction in the  $^3\Sigma_g^-$  electronic ground-state of molecular oxygen the quantum number  $J$  can take the values 0, 1 or 2 with the  $J = 2, m_J = 2$  sub-level being the most favourable for magnetic trapping or deceleration due to its ratio of mass to effective magnetic moment. Although oxygen is less massive than argon this ratio is slightly worse ( $\approx 16 \text{ Da} / \mu_B$  compared to  $\approx 13 \text{ Da} / \mu_B$  for the  $m_J = 2$  level of argon), so we would expect oxygen to have similar but somewhat worse deceleration results. The  $J = 2, m_J = 2$  level of oxygen has a linear Zeeman shift so can be easily modelled using the methods previously described for *e.g.* argon in chapter 3.

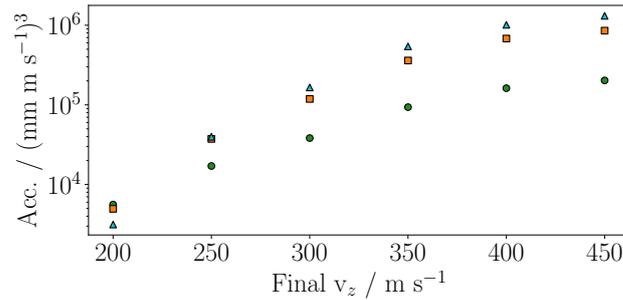


Figure 4.70: Estimated acceptance of a 20 module Durham decelerator in deceleration mode, for the  $m_J = 2$  sub-level of  $^3P_2$  argon, peak current of 350 A. The circles / squares / triangles are for no / wire / permanent magnet quadrupole.

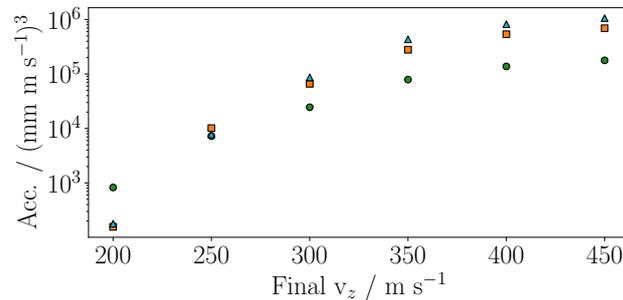


Figure 4.71: Estimated acceptance of a 20 module Durham decelerator in deceleration mode, for the  $J = 2, m_J = 2$  sub-level of oxygen, peak current of 350 A. The circles / squares / triangles are for no / wire / permanent magnet quadrupole.

The estimated acceptance of a 20 module decelerator for argon and oxygen is shown in figures 4.70 and 4.71. With an initial velocity of  $450 \text{ m s}^{-1}$  the acceptance is comparable down to a final velocity of  $\approx 300 \text{ m s}^{-1}$ , below which it drops rapidly for both species, although as expected performance is somewhat worse for oxygen. Given the similar results for the two species it seems likely that the Durham decelerator could work with such a dual-species beam, but it would not be able to decelerate it by a very useful amount unless the source was much slower. Again, we see that our performance is hampered by our low peak current at such high wave velocities.

The Tel Aviv group later reported an experiment in which molecular oxygen was seeded into a krypton beam, into which lithium was then entrained [57]. They were able to achieve an initial flow velocity of  $375 \text{ m s}^{-1}$ , then decelerate the beam and trap it in a permanent magnet quadrupole trap. We will present some preliminary simulations on proposed experiments in which we decelerate lithium with other species in chapter 5. Although they show some potential, however, these proposals rely on a much lower initial flow velocity of order  $\approx 200 \text{ m s}^{-1}$ . Another dual-species beam, this time atomic carbon and molecular oxygen, was reported by Karpov *et al.* [59]. Carbon is relatively light and in its  $^3\text{P}_1$  term the low-field-seeking  $m_J = 1$  level has a ratio of mass to moment of  $\approx 16 \text{ Da} / \mu_B$ . Choosing an initial beam velocity of  $375 \text{ m s}^{-1}$  to match that in Karpov we estimated the acceptance for a range of final velocities down to  $100 \text{ m s}^{-1}$ , which seems feasible for either a 16 or 20 module decelerator. Figure 4.72 shows the results for 20 modules.

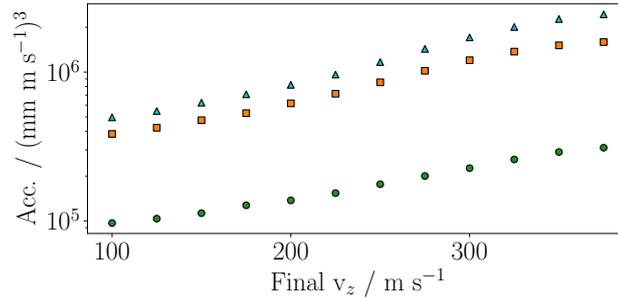


Figure 4.72: Estimated acceptance of a 20 module Durham decelerator in deceleration mode, for the  $m_J = 1$  sub-level of  $^3\text{P}_1$  carbon, peak current of 350 A. The circles / squares / triangles are for no / wire / permanent magnet quadrupole.

### 4.6.3 Basel Decelerator

The ‘Basel’ decelerator [53] is a more recent development than the experiments so far discussed. Damjanović *et al.* describe the production of a beam of hydroxyl radicals by entraining water into a supersonic expansion of xenon, with 98 % of the OH present being in the rovibrational and electronic ground state  $X^2\Pi_{3/2}$  [168]. The beam was guided at the flow velocity  $445 \text{ m s}^{-1}$  and decelerated to  $400$  and  $350 \text{ m s}^{-1}$ . Post-deceleration, the particles are detected using laser-induced fluorescence. The three experimental time-of-flight traces shown differ from those that we have seen for argon in that the bulk of the signal shifts appreciably with deceleration, rather than just the particular manipulated peaks. This is likely related to the fact that in OH there are only high-field seeking and low-field seeking states, and no ‘neutrals’ as is the case for *e.g.* argon or neon.

Damjanović report having the same problem that we faced when analysing the results of the Durham decelerator, namely that in their experiment multiple adjacent traps in the decelerator are more or less loaded, giving rise to a complicated overlapping signal that is hard to interpret. As described earlier in the context of the Paris experiment, being able to produce a sample of particles for deceleration with a small temporal or spatial width is desirable but not always feasible. Instead, the Basel group used the same method that we have described in this thesis in order to interpret their results; Monte Carlo simulation of the experiment based on multiple particle trajectories. Their simulated time-of-flight profiles are in good agreement with the data, although less so for the larger deceleration example. Their simulated phase-space plots suggest that their traps are  $\approx 6 \text{ mm}$  in length, or about a third the length of the traps in our experiment. Although spatially short samples are desirable for experimental goals such as loading a trap, this does exacerbate the problem of multiple traps being loaded unless the source is very constrained (*e.g.* short excitation pulse or a very small translational temperature, depending on the species involved).

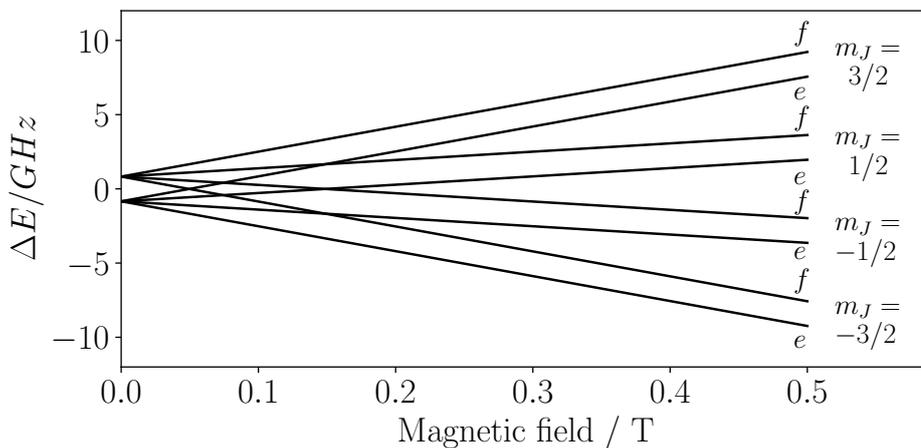


Figure 4.73: Approximate energy levels of the Zeeman splitting of OH in its rotational, vibrational and electronic ground-state  $X^2\Pi_{3/2}$ .

In order to simulate deceleration of OH we need to calculate the Zeeman splitting in an applied magnetic field  $B_z$ . The  $X^2\Pi_{3/2}$  term comprises two spectra of 8 sub-levels each, the two spectra having opposite parity [223]. At field strengths above approximately a few hundred Gauss each spectrum can be treated as four states with a linear Zeeman shift, with the two spectra separated by the ‘lambda-doubling’ energy splitting,  $\Delta_{LD}$ . The energies can be approximated by treating each

spectrum as being given by:

$$\Delta E = \mu_B(g_L\Lambda + g_S\Sigma)\frac{m_J\Omega}{J(J+1)}B_z, \quad (4.6.1)$$

Where  $m_J$  and  $\Omega$  are, respectively, the projection quantum numbers of the total angular momentum  $\mathbf{J}$  along the applied field and the molecular axis, the  $\Lambda$  and  $\Sigma$  quantum numbers represent the projection of the orbital and spin electron angular momentum along the molecular axis, and  $g_L$  and  $g_S$  are the orbital and spin g-factors. For the term in question  $\Lambda = g_L = 1$ ,  $g_S = 2$ ,  $\Sigma = 1/2$ ,  $|\Omega| = 3/2$  and  $J = 3/2$  giving  $m_J = -3/2, -1/2, 1/2, 3/2$ . This produces a set of 4 energy levels, which are duplicated and separated by  $\Delta_{LD}$ . A plot of this result is shown in figure 4.73, where  $e/f$  represent the negative / positive parity states.

Considering only the spin-stretched states labelled  $m_J = 3/2$ , we can treat them as a population with a linear Zeeman shift and effective magnetic moment  $\mu_{\text{eff}} = 6\mu_B/5$ . This was used in the static acceptance code for both the existing 4 module and proposed extended 9 module decelerator, over a range of deceleration values. The 9 module machine would be slightly longer than the Basel decelerator. The results are shown in figures 4.74 and 4.75.

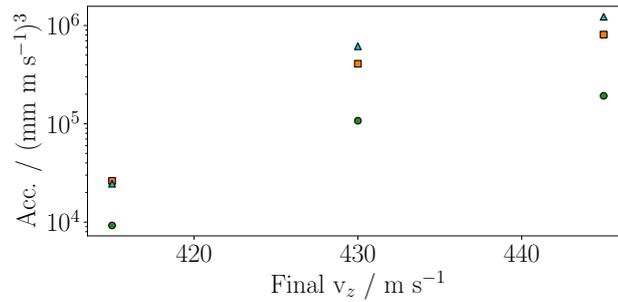


Figure 4.74: Estimated acceptance of the Durham decelerator in deceleration mode, for the  $m_J = 3/2$  sub-levels of  $X^2\Pi_{3/2}$  hydroxyl. The circles / squares / triangles are for no / wire / permanent magnet quadrupole. Decelerator operating at 350 A.

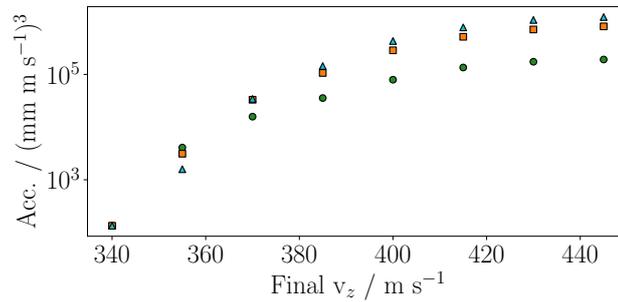


Figure 4.75: As per figure 4.74, but for the proposed 9 module decelerator.

The acceptance figures suggest that the existing decelerator should be able to successfully guide an OH beam at  $445 \text{ m s}^{-1}$  and decelerate a portion of it to perhaps  $430 \text{ m s}^{-1}$ , but the acceptance drops rapidly for harder deceleration. The proposed 9 module extension has more promising acceptance down to a final velocity of  $400 \text{ m s}^{-1}$  and perhaps further. To test this deceleration sequences were simulated to match the initial and final velocities reported in Damjanović [168], with a peak decelerator current of  $350 \text{ A}$  and wire quadrupole at  $700 \text{ A}$ . An initial OH distribution was generated using equation 3.2.2, with  $v_T = 445 \text{ m s}^{-1}$  and  $\sigma_{T,z} = 50 \text{ m s}^{-1}$  (as reported in Damjanović) but other source parameters typical for our experiment. Although rather artificial this should at least give us some insight into the relative performance of the two decelerators.

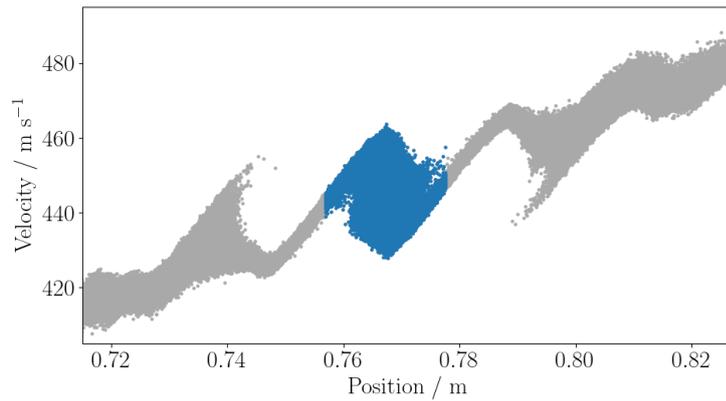


Figure 4.76: Longitudinal phase-space plot of simulated guiding of  $m_J = 3/2$  sub-levels of  $X^2\Pi_{3/2}$  OH at  $445 \text{ m s}^{-1}$  in the 4 module Durham decelerator, peak current of  $350 \text{ A}$ , wire quadrupole at  $700 \text{ A}$ . Particles within the target trap maxima are in blue.

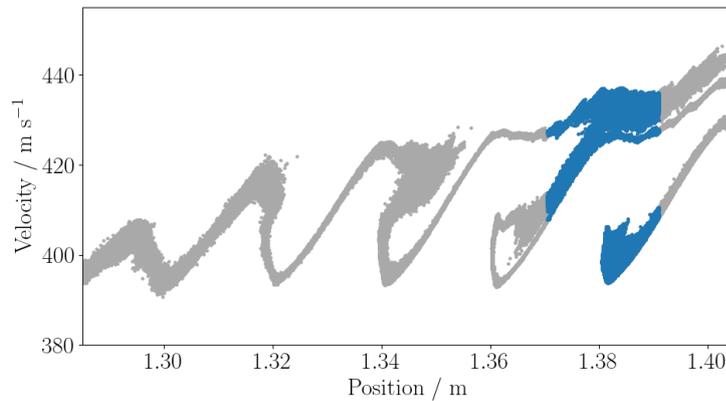


Figure 4.77: Longitudinal phase-space plot of simulated deceleration of  $m_J = 3/2$  sub-levels of  $X^2\Pi_{3/2}$  OH from  $445$  to  $400 \text{ m s}^{-1}$  in the proposed 9 module Durham decelerator, peak current of  $350 \text{ A}$ , wire quadrupole at  $700 \text{ A}$ . Particles within the target trap maxima are in blue.

Good results were obtained for the 4 module decelerator for guiding at  $445 \text{ m s}^{-1}$ , figure 4.76 shows a longitudinal phase-space plot of the particles at the end of the guiding sequence. Deceleration from  $445$  to  $400 \text{ m s}^{-1}$  gave a weak result, there was evidence of deceleration but the separatrix is small for these parameters. The proposed 9 module decelerator performed well for both guiding and deceleration to  $400 \text{ m s}^{-1}$ , figures 4.77 and 4.78 show phase-space and velocity plots of the decelerated atoms. There are faster non-phase-stable particles exiting the trap, but the central bunch has been clearly decelerated by the target amount. At a final velocity of  $350 \text{ m s}^{-1}$  we again saw a weak performance compared to the Basel decelerator.

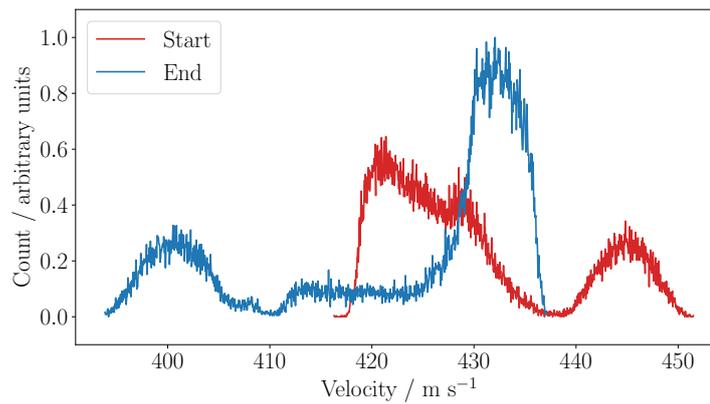


Figure 4.78: As per figure 4.77, longitudinal velocities of particles loaded into the target trap (red) before / (blue) after deceleration.

Figures 4.76 and 4.77 can be compared to figures 11a and 11b in Damjanović [168]. The Basel and Durham phase-stable regions appear to be comparable in size in velocity space for guiding, although for deceleration the Basel result is perhaps twice as large. The greater longitudinal size of the Durham traps is clear. It is unclear what the temporal width of the Basel source is, but it appears to be larger than the width of our source judging by the number of filled traps, even allowing for the smaller trap size.

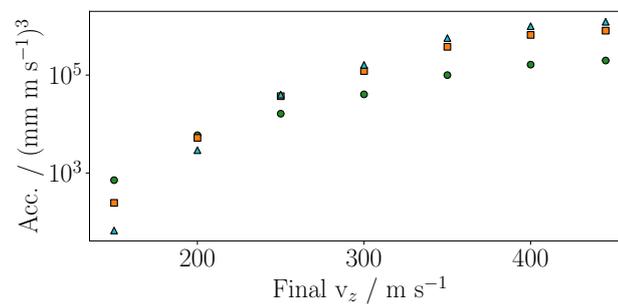


Figure 4.79: As per figure 4.74, but for a 20 module decelerator.

Figure 4.79 shows estimated acceptance for OH in a 20 module Durham decelerator. This machine would be significantly longer than the Basel decelerator and we have considered deceleration of OH to much lower velocities than they have reported; however it seems that with the initial conditions reported by Damjanović the Durham design is less capable than the Basel for decelerating OH.

#### 4.6.4 Conclusion

To summarise this section; we have used our simulation codes to try and compare our decelerator with several other designs. The Durham decelerator is easily programmed with a desired target travelling-wave profile using the technique of pulse-width modulation. Issues relating to oscillations of the trap centre and depth caused by the synthesis of a ‘moving’ magnetic trap are not unique to our design but do not seem to be a serious impediment to performance.

The decelerator can provide strong longitudinal confinement, although this is dependent on the required wave velocity. It is not well suited to dealing with beams with a high initial flow velocity, especially in its current form. Based on our simulation work it shows promise for handling certain species, *e.g.* Li or H, especially if it is extended as we have previously proposed [2], however the feasibility of such possible future work strongly depends on the nature of the beam source employed.

A significant issue with the Durham decelerator design is that of transverse focussing, as considered at length in section 4.5.2. The Tel Aviv design does not have the same problems with lateral confinement due to the rotationally symmetric nature of its fields, although periodic variation in radial trap depth as the beam passes between coil pairs can cause coupling of longitudinal and transverse velocities, leading to instability. Ji *et al.* claim to have designed a similar machine that addresses this problem through an amended method of applying the time-dependent currents to the coil phases [55]. The Basel design, as discussed, has a similar field profile to the Durham decelerator except that it rotates around the beam axis. With the correct choice of travelling-wave parameters this seems to reduce transverse loss.

Compared to the Paris decelerator upon which it was based we believe we have improved upon its performance somewhat, but certainly the decelerator in its current state is of limited potential compared to the other experiments now in use. A future development of this experiment would need to extend the number of decelerator modules and consider carefully the issue of focussing; whether to proceed with the planned permanent magnet quadrupole or a revised version of it, or return to some form of wire quadrupole. Probably the most serious impediment to its performance is the fact that it can only achieve close to its maximum current when the travelling-wave velocity is low, making it best suited for beams that are already slow in the laboratory frame. This will be further discussed in section 5.1.

# Chapter 5

## Outlook

Our travelling-wave Zeeman decelerator was initially conceived as part of a project with the goal of obtaining trapped samples of ultracold molecules. The physical design of the experiment is discussed in chapter 2, and chapter 3 analysed some of the data we have taken when characterising the machine with metastable argon. Although this data is sub-optimal in terms of the experimental parameters that were chosen, we have demonstrated that the experiment operates in the way that we claim, and we are able to utilise the suite of software tools that we have developed to interpret our data and make predictions about our expected results for different parameters, or even physical alterations to the apparatus. Additionally, the decelerator on its own is not enough; in order to achieve our research goals, some form of trap must be built and added to catch the output of the decelerator. These topics will be the subject of this chapter, which will consider the short- and long-term outlook for the experiment, and describe the results of some preliminary design and simulation work that we have undertaken.

### 5.1 Extending the Decelerator

McArd [2] characterised the effects of deceleration on the moving traps in the travelling wave in two ways; firstly, by estimating the phase-space acceptance of the traps for a given species via the method described in section 4.3. Secondly, he presented the trapping potential of the decelerator wave, modified by treating the deceleration as a pseudoforce that acts to alter the shape and depth of the walls of the trap. This view is very illustrative, but in the present work we have focussed on ways of estimating the acceptance and displaying the results as phase-space plots. To achieve this we have expanded on the ‘static field’ method previously used, and developed the fully dynamic model described in section 4.4.

Both the work of McArd and the work described in this thesis have shown that the decelerator in its current 4 module long form is too short to decelerate metastable argon to an effective standstill, given a beam with initial conditions achievable with our current source. The greater the magnitude of the deceleration the smaller the volume of the phase-stable region in the decelerator traps, and hence the smaller the acceptance. For species with a more favourable ratio of magnetic moment to mass the machine could effectively remove a greater fraction of the forward velocity of the beam, for example some exploratory simulations of atomic hydrogen suggest that it could be significantly decelerated even with the 4 module machine, dependent on the initial flow velocity and temperature - this can be seen in the estimated acceptance shown in figure 4.40 in section 4.4, for which the initial beam velocity was  $350 \text{ m s}^{-1}$ . (More detail will be given in section 5.3.4).

McArd discussed a proposed sympathetic cooling experiment in which lithium and calcium monohydride were co-decelerated and loaded into a hybrid magnetic trap / magneto-optical trap (MT-MOT), further details of which will be given later in the chapter. It was estimated that a 9 module decelera-

tor would be required in order to decelerate the particles to a low enough velocity that they could be loaded into the trap, whilst achieving adequate acceptance. Some predictions related to deceleration of argon with this extended machine were given in chapter 4 for the purposes of comparison with our existing argon data. As discussed in section 4.6, however, a 9 or indeed 16 or 20 module decelerator is still not able to match the performance of some other travelling-wave decelerators, dependent on the species and initial beam conditions. A longer decelerator coupled with a slower source could potentially allow us to handle heavier species, for which the deeper traps we can generate at low frequency would be an advantage.

Some caveats must be made regarding a longer decelerator. In the speculative simulations shown here we have included a wire quadrupole, which we previously had installed. Extending this to 9 modules or longer would be challenging, due to the technical difficulty of running such high currents through a long wire. The proposed permanent magnet quadrupole has no such limitation as to length, however as discussed in section 4.5.2 there are possible problems with the current design at low beam velocities. More work remains to be done to analyse this. If a wire quadrupole was to be used again, then care should be taken regarding its polarity. In previous chapters we showed that the combination of wire quadrupole and decelerator field causes alternate traps to be strengthened or weakened, to which end we typically aimed to load the second full trap of the travelling-wave. However, particles that are in the second trap experience some time at the end of the deceleration sequence during which they have left the field but it is still active behind them. For relatively fast beams of heavy species such as argon this is not a significant problem but for light / slow species such as hydrogen the decelerated packet of atoms gets a significant kick back to higher mean velocity in the time it takes the field to collapse. Unless the pulse-width modulation sequence is revised, the best approach would be to load the third trap in the wave, but have the wire quadrupole installed with reversed polarity with regard to the decelerator coils.

An alternative to simply extending the length of the Durham decelerator is to combine it with a conventional Zeeman decelerator, which could provide a beam with a lower mean velocity to input into the travelling-wave machine. This would allow us to work with a greater variety of ‘fast’ sources but then utilise the deep trap depths available when operating at lower frequencies and thus higher currents. A similar approach was described by Quintero-Pérez *et al.* [39]; a conventional Stark decelerator was used to feed beams of  $\text{NH}_3$  and  $\text{ND}_3$  into a travelling-wave Stark decelerator, combining the best features of both devices. This could be explored in further simulation work without difficulty; we have previously used our codes to model a conventional Zeeman decelerator for a proposed experiment [224] in which the ‘photostop’ technique [125] was used to load a decelerator.

In chapter 4 we have already considered the efficacy of an extension to the existing decelerator in the case of metastable argon as well as some other species used by other groups. In this chapter we will make further predictions on other species of interest to us.

## 5.2 Magnetic Trap / Magneto-Optical Trap

Magneto-optical traps or ‘MOTs’ are generally built with a pair of coils in anti-Helmholtz configuration to provide a quadrupole magnetic field at the centre of the trap. These field gradients give rise to position-dependence of the laser scattering effects, creating forces that act to confine the trapped particles. For sympathetic cooling of species that cannot be cooled or trapped with lasers we require additional magnetic trapping (see section 3.4). The proposed ‘magnetic trap / magneto-optical trap’ or ‘MT-MOT’ is shown in figure 5.1; it comprises 6 toroidal magnets to give a trapping potential with steep sides to provide strong confinement of paramagnetic species, but with shallow gradients at the centre that should allow laser-cooling.

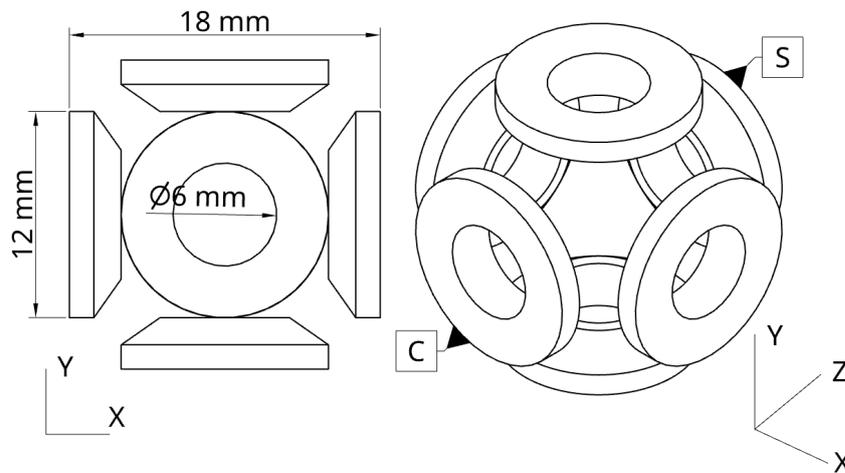


Figure 5.1: The magnetic trap / magneto-optical trap design, comprising 6 toroidal magnets. The face labelled ‘c’ contains a ‘cancellation’ coil, when active this nullifies the field due to the front magnet and effectively ‘opens’ the trap. The face labelled ‘s’ contains a ‘slower’ coil, which is switched on during loading to provide a steeper stopping potential.

Some analysis of this trap geometry was carried out by D. Nohlman [225] with further work done by the current author [151]. The detailed design is in McArd [2]. The MT-MOT is intended to be mounted at the end of the decelerator; the face that abuts the decelerator contains a coil calculated to cancel out the magnetic field due to the front magnet. This ‘opens’ the trap and allows the decelerated beam to be loaded. An additional coil is mounted around the opposite, rear magnet; this is designed with current polarity such that it reinforces the magnet to create a much steeper final stopping potential. The intended mode of operation is that the switching times of the two coils are set so as to allow a packet of particles to enter the trap and be virtually stopped at the rear side, the slower coil is then switched off before the particles roll back down its potential gradient and regain their lost kinetic energy - it thus works in the same way as an individual coil in a convention Zeeman decelerator. The cancellation coil is then switched off before the particles can rebound back and escape the trap. For a detailed description of the shape of the trapping fields and earlier work on calculating optimised coil switching sequences see Walker [151]. Ideally, the MT-MOT is then loaded with a population of two species, both magnetically trapped; one that can be laser-cooled, and one that cannot. Taking the beam loading direction to be along the  $z$ -axis, counter-propagating lasers are mounted along  $x$  and  $y$ , as is conventional for magneto-optical traps. Where this design differs is that we cannot easily place a laser along  $z$ , and so it is proposed to add the required  $z$ -component along one of the trap diagonals.

### 5.3 Sympathetic Cooling

The Zeeman decelerator described in this thesis was originally conceived as being an enabler for a molecular quantum simulation experiment, by providing an improved mechanism for the deceleration of paramagnetic molecules. A key factor in this goal and an interesting outcome in its own right is the necessity to understand and control the dynamics of cold molecular collisions. Sympathetic cooling offers a potential alternative route to ultracold molecules, compared to other well-known techniques such as magneto-association. In simple terms, it consists of co-trapping the molecular species of interest with some pre-cooled atomic species. The cooled atoms then collide with the molecules and thermalise, removing energy from the molecules. The atoms may be laser-cooled initially. Evaporative cooling may be applied after the two species are co-located so that the atoms extract energy from the molecules and are then removed from the trap. The feasibility of this process is strongly dependent on the properties of the system involved; the particles must be suitable for magnetic trapping, so that they can be confined whilst the laser and collisional cooling processes occur. (For species / states that cannot be magnetically trapped, confinement using ac [135] or microwave [226] trapping may be an alternative approach, but only magnetic trapping is considered here).

Inelastic collisions can cause the atoms or molecules to flip to a high-field-seeking state and be ejected from the trap, and a common ‘rule-of-thumb’ figure given is that the ratio of elastic to inelastic collisions should be at least  $\approx 100 : 1$  [227]. It is also desirable for the molecular species to be described by Hund’s case b, *i.e.* the electron spin coupling to the molecular axis is weak, since in this case the inelastic cross-section is smaller [228]. In collisions of a  $^2\Sigma^+$  diatomic molecule with an atom the expected rate of loss of trapped particles due to inelastic collisions is dependent on the interaction of the spin and rotation of the molecule [229]. This is reduced for the rotational ground-state  $N = 0$  state, in which the spin-flip transitions depend on coupling to the  $N > 0$  states [230]. Studies also show that it is advantageous if the atomic species is light [231]; this gives rise to large centrifugal barriers which help suppress low energy losses. This has been tested in, for example, collisions between NH and H [232] or Li [233]. In the next section we consider interactions of calcium monohydride, CaH ( $^2\Sigma^+$ ) and lithium, Li ( $^2S$ ).

### 5.3.1 Calcium Monohydride plus Lithium

An interesting system for possible sympathetic cooling is that of calcium monohydride, CaH ( $^2\Sigma^+$ ) and lithium, Li ( $^2S$ ). This system has interactions that offer great scope for control in the cold and ultracold regimes. The low reduced mass of the system makes it a good choice for energy exchange through collision. In a magnetic trap the net angular momenta of the particles must be projected in the same direction relative to the space quantisation axis, so that the exothermic singlet reaction surface is not present. The triplet surface is endothermic, so that at low particle energies it is also inhibited.

Even with reactions suppressed, inelastic collisions are of concern - as previously mentioned, a ratio of elastic-to-inelastic collisions of at least 100 to 1 should exist for sympathetic cooling. If optically pumped into their maximally (electron) spin-stretched states, theory suggests that the cross-sections for the CaH + Li system are very favourable for this [69]. These estimated cross-sections, however, did not include hyperfine effects - it remains to be tested whether the presence of states with opposing nuclear spins may interfere with the suppression of inelastic collisions. Feshbach resonances have been found to be important in inelastic collision rates in other systems [234]. The estimated cross-sections are also strongly affected by the accuracy of the *ab initio* calculated potential surfaces [73]. Since we originally considered this system a new *ab initio* potential energy surface has been calculated [235], and this remains a system of interest.

Additionally, spin-orbit interactions may lead to non-adiabatic transitions (also known as ‘Majorana transitions’ or ‘spin-flips’) to the singlet surface. Theory suggests that these transitions should be negligible for the CaH / Li system [71]. Unfavourable predictions were made for the proposed MT-MOT field geometry in work by another group [225], however, this work relied on assumptions that were strictly only applicable in quadrupole traps. These tests were subsequently expanded upon by the current author [151] with much more encouraging results, although further work remains to be done. Our group has also studied the possible addition of a time-orbiting-potential or ‘TOP’ field to the MT-MOT [236]; TOP traps work by creating an effectively rotating bias field in addition to the trapping fields. A prototype was constructed and the fields characterised, with simulations of Li in an MT-MOT plus TOP system carried using the codes described in this thesis. Some preliminary assessment of the probability of Majorana transitions in the travelling-wave decelerator is discussed in section 5.5.

### 5.3.2 Co-Deceleration and Trapping of CaH and Li

As mentioned in section 5.2 we previously simulated the loading of the MT-MOT [151] with a dual-species beam of Li and CaH; the initial conditions were given by the output of earlier decelerator models [2]. As described in chapters 3 and 4 more detailed codes for the decelerator have been developed; the following sections give results of their application to Li and CaH. The codes can be configured to model beams composed of multiple species with different distributions of initial positions and velocities, however for clarity we will show the results for Li and CaH separately.

The ground state term of lithium is  $^2S_{1/2}$ , with electron spin 1/2, zero orbital angular momentum and nuclear spin 3/2. Using the methods outlined in section 3.4 we can calculate the splitting of the

energy levels of this atom in an external magnetic field, as seen in the Breit-Rabi diagram in figure 5.2. Hyperfine interactions give rise to two manifolds, which further separate under Zeeman effects to give four low-field-seeking and four high-field-seeking sub-levels. The desired spin-stretched  $F = 2, m_F = 2$  sub-level is highlighted in red; this has an effectively linear energy shift in an external field, and an effective magnetic moment of  $1 \mu_B$ . A sample of Li containing all of these sub-levels could be modelled using the eigenenergy interpolation code described in section 3.4, but for the purposes of this simulation we have considered a population composed of only atoms in the upper state. Given its linear Zeeman shift we can determine the force arising from interactions with the decelerator or trap fields simply from the gradient of the field at the required point.

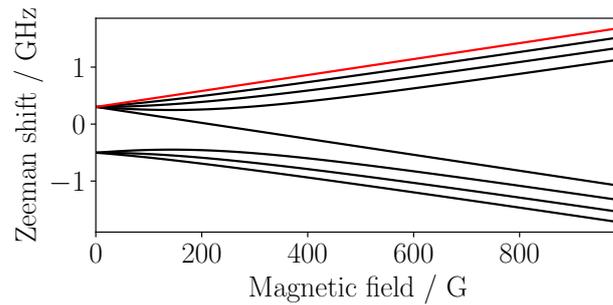


Figure 5.2: Energy splitting of the sublevels of the  $^2S_{1/2}$  state of lithium in an applied magnetic field. The maximally spin-stretched low-field-seeking  $F = 2, m_F = 2$  sublevel is shown in red.

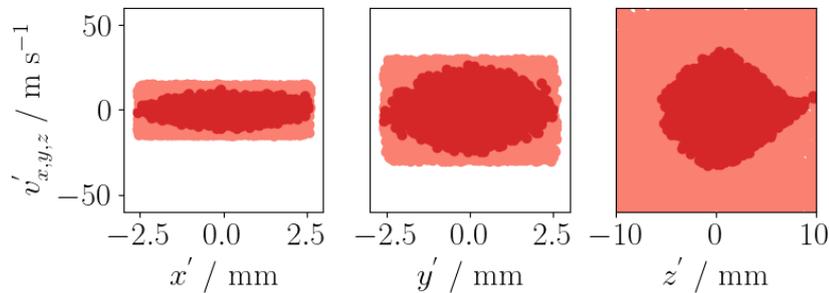


Figure 5.3: Acceptance of the  $F = 2, m_F = 2$  sub-level of  $^2S_{1/2}$  lithium in the Zeeman decelerator, operating at 1000 A peak current, decelerating from 200 to  $20 \text{ m s}^{-1}$ . Estimated via the static simulation method described in section 4.3. Pink, initial atom distribution, red, those remaining in the decelerator field after 2 ms. Coordinates are in a reference frame centred on the trap.

In this example we set the (9-module) decelerator to have initial / final trap velocities of  $200 / 20 \text{ m s}^{-1}$  with peak current of 1000 A, and the wire quadrupole active. Using the techniques described in section 4.3 the acceptance of a single trap in the decelerator field was estimated to be  $2.6 \times 10^6 (\text{mm m s}^{-1})^3$ , as illustrated in figure 5.3. For comparison, running the simulation with no quadrupole gave a value of  $5.9 \times 10^5 (\text{mm m s}^{-1})^3$ , and the permanent magnet quadrupole gave  $1.8 \times 10^6 (\text{mm m s}^{-1})^3$ . These numbers agree with those found by McArd [2] to within  $\pm \approx 10\%$ .

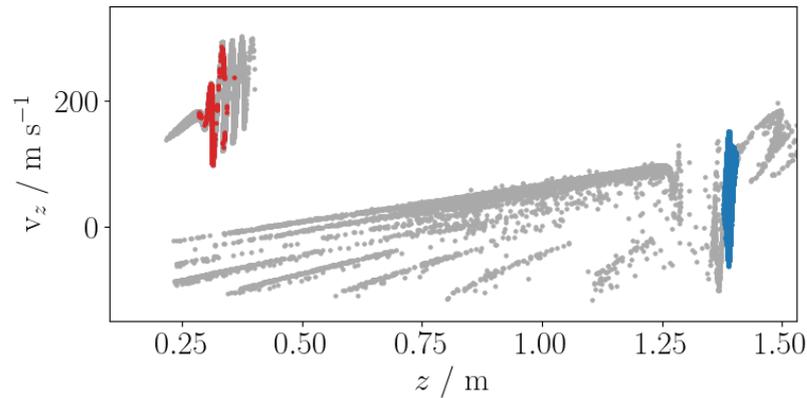


Figure 5.4: Deceleration of the  $F = 2, m_F = 2$  sub-level of  $^2S_{1/2}$  lithium, decelerator operating from  $200$  to  $20 \text{ m s}^{-1}$  with peak current of  $1000 \text{ A}$ . Upper left, the atoms just as the second trap forms, lower / right, the atoms just as the trap collapses. The atoms located within the bounds of the trap at the end are in blue, their initial coordinates are in red.

The deceleration of an initial population of Li atoms with flow velocity  $v_z = 200 \text{ m s}^{-1}$  and longitudinal spread of velocities  $\sigma_{v_z} = 13 \text{ m s}^{-1}$  was modelled. This corresponds to a translational temperature  $\approx 0.14 \text{ K}$ , chosen for comparison with the parameters tested in McArd [2]. The decelerator delay was determined such that the second trap in the wave would be optimally overlaid by the Li beam at the instant it formed. Figure 5.4 shows the longitudinal phase-space coordinates of the beam at the times that trap 2 forms and collapses, at  $1.5$  and  $11.2 \text{ ms}$ , respectively. Compared to similar plots for argon in section 3.10 several differences are immediately apparent. For these conservatively-estimated initial conditions, namely a slow and internally cold beam, the beam is much less spread out in phase space than the hotter, faster argon. The decelerator field also has much deeper traps in this simulation, given the higher current, and the atoms have interacted with them for longer by the time the trap is formed than was previously the case. Coupled with the lighter mass of Li than Ar, therefore, the atoms have already been significantly bunched by the time the first few traps form. For this longer 9-module decelerator gentler deceleration is possible; along with the deep traps this gives a large phase-stable region, and it can be seen in the bunched group of atoms at the right that the bulk of the forward velocity has been shifted down. Also different from the Ar simulations is the presence at the end of a significant population of atoms with negative velocities in the  $z$ -direction, these have effectively ‘dropped out’ of the back of the decelerator wave and been pushed back towards the source. Figure 5.5 shows the initial (red) and final (blue) velocities in  $z$  of only the atoms that are located in trap 2 as the field collapses, the mean values of the red / blue traces are  $181.7 / 44.5 \text{ m s}^{-1}$ .

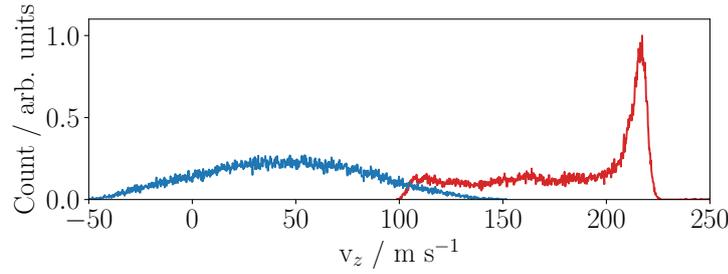


Figure 5.5: Deceleration of the  $F = 2, m_F = 2$  sub-level of  ${}^2S_{1/2}$  lithium, decelerator operating from 200 to 20  $\text{m s}^{-1}$  with peak current of 1000 A. Only the atoms remaining in the second trap of the decelerator wave at the end are considered. Red, their longitudinal velocities just as the trap forms, blue, as it collapses.

Section 2.3.1 discussed the Zeeman shift of atomic energy levels in an external field, with methods of calculating these shifts for various species described in section 3.4. Molecules also undergo energy shifts in applied fields, but the situation is greatly complicated by the presence of rotational and vibrational terms in the overall Hamiltonian. In some cases a simplified molecular Hamiltonian may give adequate results [237]. Calcium monohydride has electronic ground state  $X^2\Sigma^+$ , with no electronic orbital angular momentum and a nuclear spin of 1/2. Similarly to the atomic case we can describe the total angular momentum as  $\mathbf{F} \equiv \mathbf{J} + \mathbf{I}$  where  $\mathbf{J}$  and  $\mathbf{I}$  are the electronic and nuclear terms respectively, but in this case  $\mathbf{J} = \mathbf{N} + \mathbf{S}$ , with the term  $\mathbf{N}$  arising from the coupling of  $\mathbf{L}$  with the nuclear rotational angular momentum  $\mathbf{R}$ . The simplified Hamiltonian therefore contains contributions from the electron spin, molecular rotation, Zeeman effects and corrections to centrifugal distortion, plus coupling of these terms. The energy levels are shown in figure 5.6, and the low-field-seeking  $m_J = 1/2$  sub-level used in the deceleration simulations is marked in red.

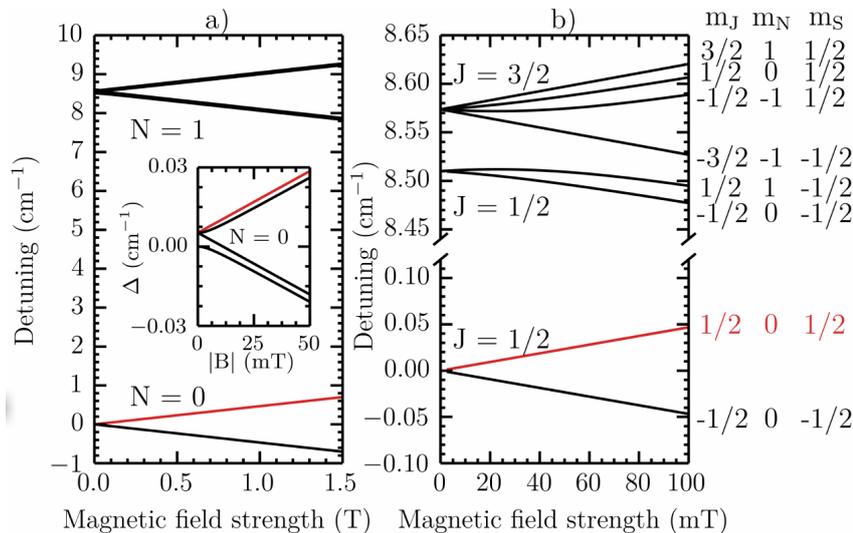


Figure 5.6: Energy splitting of the electronic ground state  $X^2\Sigma^+$  of calcium monohydride in an applied magnetic field. a): The rotational states  $N = 0$  and 1. Inset, hyperfine splitting of  $N = 0$ . b): Zero-field splitting of the rotational states and Zeeman-shifting of the sublevels described by the  $m_J$  quantum number. The  $m_J = 1/2$  state used in the decelerator simulations is shown in red. Image reproduced from McArd [2].

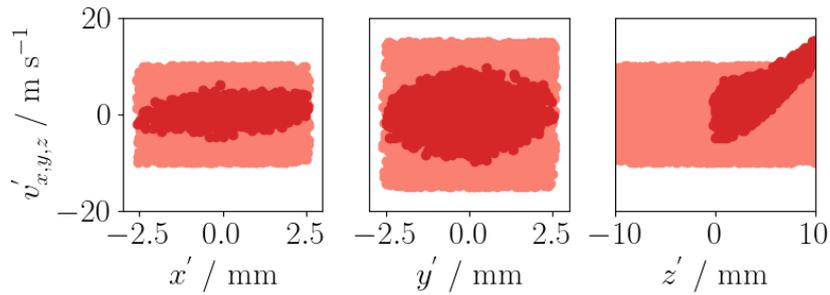


Figure 5.7: Acceptance of calcium monohydride in the Zeeman decelerator, operating at 1000 A peak current, decelerating from 200 to 20  $\text{m s}^{-1}$ . Estimated via the static simulation method described in section 4.3. Pink, initial atom distribution, red, those remaining in the decelerator field after 2 ms. Coordinates are in a frame centred on the moving trap.

The acceptance of CaH in the decelerator trap was estimated to be  $3.1 \times 10^3 (\text{mm m s}^{-1})^3$ , as shown in figure 5.7. This is likely to be an upper boundary on the actual acceptance, since it can be seen that some phase-unstable molecules are still escaping. However, simulating the deceleration of this species using the same parameters as given for Li gave promising results, as seen in figures 5.8 and 5.9. We have chosen an initial translational temperature of 20 mK for the beam; as was the case for Li this value was chosen for comparison with the earlier work of McArd [2]. This (admittedly conservatively low) temperature gives a small spread in position by the time the molecules are being loaded into the trap. Given the effective deceleration shown in figure 5.9, in which the mean longitudinal velocity shifted from 199.0 to 27.1  $\text{m s}^{-1}$ , it seems likely that somewhat hotter beams could be handled.

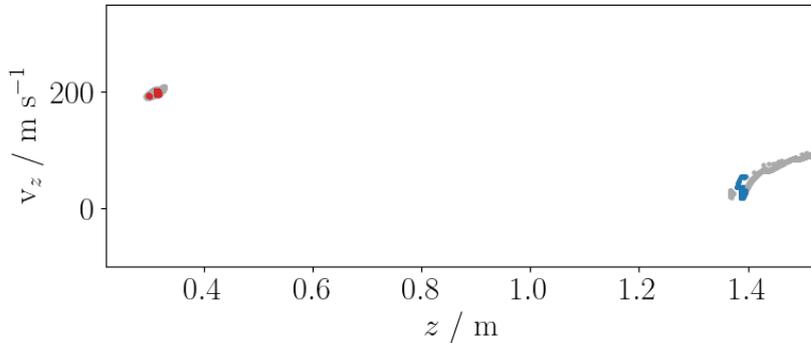


Figure 5.8: Deceleration of calcium monohydride, decelerator operating from 200 to 20  $\text{m s}^{-1}$  with peak current of 1000 A. Upper left, the molecules just as the second trap forms, lower right, the molecules just as the trap collapses. The atoms located within the bounds of the trap at the end are in blue, their initial coordinates are in red.

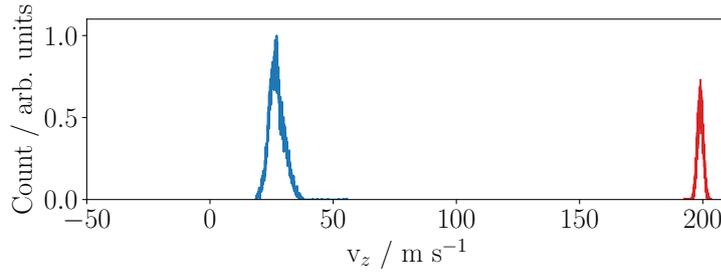


Figure 5.9: Deceleration of calcium monohydride, decelerator operating from 200 to 20  $\text{m s}^{-1}$  with peak current of 1000 A. Only the molecules remaining in the second trap of the decelerator wave at the end are considered. Red, their longitudinal velocities just as the trap forms, blue, as it collapses.

Utilising the method described in section 4.3 we can estimate the acceptance of the MT-MOT for both Li and CaH, obtaining values of  $5.2 \times 10^7 (\text{mm m s}^{-1})^3$  and  $5.3 \times 10^6 (\text{mm m s}^{-1})^3$  respectively. Figures 5.10 and 5.11 show the initial and final distributions of position and velocity for the particles along the axes. Comparing the phase-stable regions with those in figures 5.3 and 5.7 confirm that the acceptance of the trap is larger than the emittance of the decelerator.

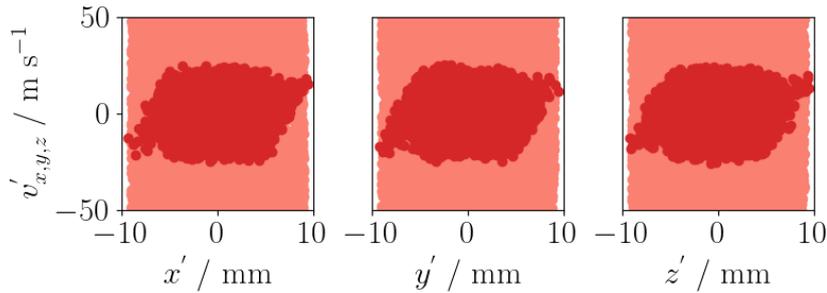


Figure 5.10: Acceptance of the  $F = 2, m_F = 2$  sub-level of  $^2\text{S}_{1/2}$  lithium in the MT-MOT. Pink, initial atom distribution, red, those remaining in the trap field after 2 ms.

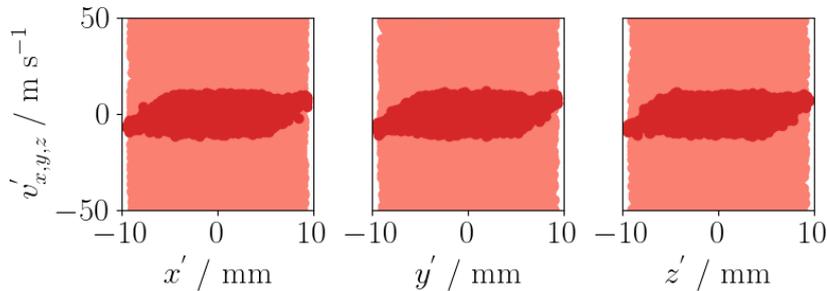


Figure 5.11: Acceptance of calcium monohydride in the MT-MOT. Pink, initial atom distribution, red, those remaining in the trap field after 2 ms.

Together, these preliminary results suggest that the combination of travelling-wave decelerator and MT-MOT has potential for exploring low-temperature chemistry with the CaH + Li system. Future work could refine the optimisation of trap-loading for these species, and incorporate the actual laser- and collisional-cooling processes.

### 5.3.3 Hydrogen plus Lithium

Changing focus from controlled low-temperature chemistry and cold molecules; there has been significant recent interest in the search for new physics. One proposed method of achieving the required spectroscopic precision to test for deviations from the Standard Model is to use light atoms for which the atomic wavefunction can be calculated very accurately. For example; measuring the Rydberg constant for hydrogen across Rydberg states from  $n = 2$  to  $n = 30$ , with a corresponding range of wavefunction radii up to  $r \approx 30 a_0$ , would allow constraints on the masses of new bosons in the eV to keV regime to be calculated. Reaching the ultracold regime with hydrogen is difficult; it cannot be laser-cooled due to limitations in current technology [238]. It has been evaporatively cooled but this process is very slow due to the small S-wave scattering length [239]. It has been suggested that sympathetic cooling of hydrogen by lithium is feasible [240]. We have therefore considered using our travelling-wave Zeeman decelerator to produce a dual-species beam of atomic hydrogen and lithium, with the aim of loading it into the MT-MOT. We have carried out simulations of the deceleration of atomic hydrogen, and estimated the acceptance of the decelerator and trap for this species; similar results for lithium were given in section 5.3.2. Additional work was done on simulating the MT-MOT by O. Turnbull [241], who began combining the interpolation, magnetic trapping and numerical integration codes described in chapter 3, with laser-scattering code based on that of R. Hanley *et al.* [242]. This was greatly developed by S. Booshanam [189], who produced some preliminary simulations of laser-cooling of Li and collisional cooling between Li and H in the MT-MOT.

### 5.3.4 Co-Deceleration and Trapping of H and Li

The energy levels of  $^2S_{1/2}$  hydrogen are shown in figure 3.12. If it is required to handle all four sub-levels and their motion in a magnetic field we can use the eigenenergy interpolation scheme described in section 3.4; for the purposes of the deceleration simulations shown here we can instead consider only the  $F = 1, m_F = 1$  state, as we did for lithium in section 5.3.2. The force on the atoms is then, once again, easily related to the gradient of the magnetic field.

Figure 4.40 in section 4.4 shows the estimated acceptance of the spin-stretched  $m_F = 1$  level of ground-state hydrogen for deceleration from an initial velocity of  $350 \text{ m s}^{-1}$ , chosen to be comparable to the typical velocities we see when producing a supersonic expansion of argon. Figure 5.12 shows a simulation of hydrogen deceleration from  $350$  to  $100 \text{ m s}^{-1}$  with peak current of  $400 \text{ A}$  in the existing 4-module decelerator, with a view of the velocity distributions of the atoms pre- and post-deceleration in figure 5.12. This represents a reduction of  $\approx 98 \%$  of the kinetic energy of the atoms, which are well bunched in several following traps, with the faster non-phase-stable atoms escaping above.

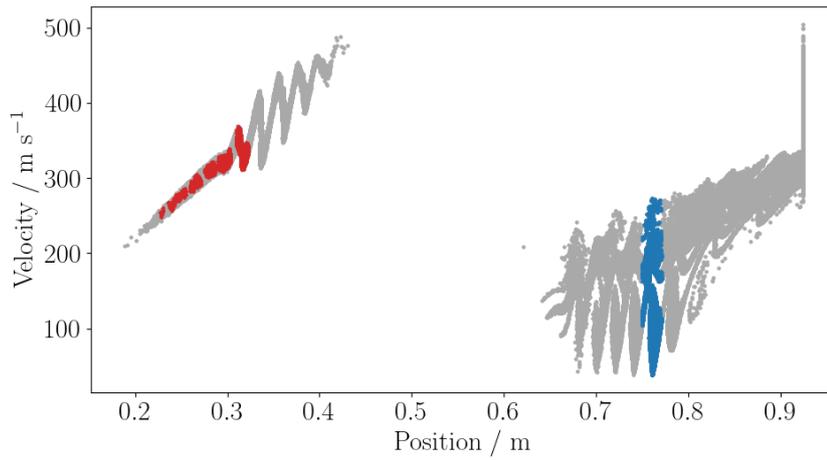


Figure 5.12: Simulated deceleration of the  $F = 1, m_F = 1$  sub-level of  $^2S_{1/2}$  hydrogen, 4-module decelerator operating from  $350$  to  $100 \text{ m s}^{-1}$  with peak current of  $400 \text{ A}$ . Upper left, the atoms just as the second trap forms, lower / right, the atoms just as the trap collapses. The atoms located within the bounds of the trap at the end are in blue, their initial coordinates are in red.

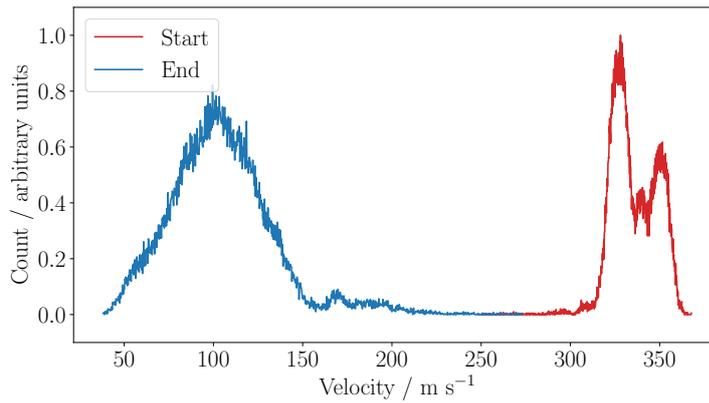


Figure 5.13: As per figure 5.12, but comparing the longitudinal velocities of the atoms remaining in trap 2 as it collapses (blue) with their velocities as the trap formed (red).

Once again, we consider a slower beam, decelerated from  $200$  to  $20 \text{ m s}^{-1}$  in the proposed 9-module decelerator. The estimated acceptance of hydrogen in one trap of the decelerator wave is  $3.9 \times 10^7 (\text{mm m s}^{-1})^3$ , with the results of the simulation code shown in figure 5.14. Figure 5.15 shows the result of simulating the deceleration of hydrogen, with the (9-module) decelerator parameters the same as in the previous sections;  $1000 \text{ A}$  peak in the deceleration coils, wire quadrupole active. The source is created with a mean flow velocity  $v_z$  of  $200 \text{ m s}^{-1}$  and a standard deviation  $\sigma_{v_z} = 28 \text{ m s}^{-1}$  corresponding to a temperature of  $\approx 95 \text{ mK}$ . Very similar structure to the lithium simulation shown in figure 5.4 is seen; there is a strongly spatially-bunched cluster of atoms emerging from the decelerator trap as it collapses, with a significant spread of velocities extending to negative values. There is also a secondary cluster with a very similar shape trailing it, and significant backwards scattering of particles that been lost out of the rear of the decelerator wave and are moving back towards the origin. The pre- and post-deceleration longitudinal velocities of just the atoms in this trap are shown in figure 5.16; these atoms make up  $\approx 2.7 \%$  of the initial sample, and their mean velocity in  $z$  has decreased from  $180.9$  to  $44.1 \text{ m s}^{-1}$ .

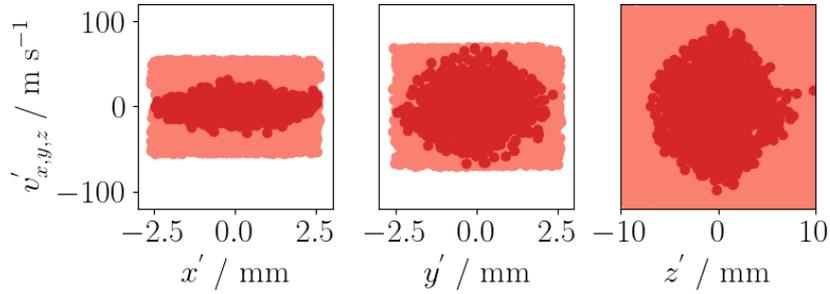


Figure 5.14: Acceptance of the  $F = 1, m_F = 1$  sub-level of  ${}^2S_{1/2}$  hydrogen in the 9-module Zeeman decelerator, operating at 1000 A peak current, decelerating from 200 to 20  $\text{m s}^{-1}$ . Estimated via the static simulation method described in section 4.3. Pink, initial atom distribution, red, those remaining in the decelerator field after 2 ms. Coordinates are in a reference frame centred on the trap.

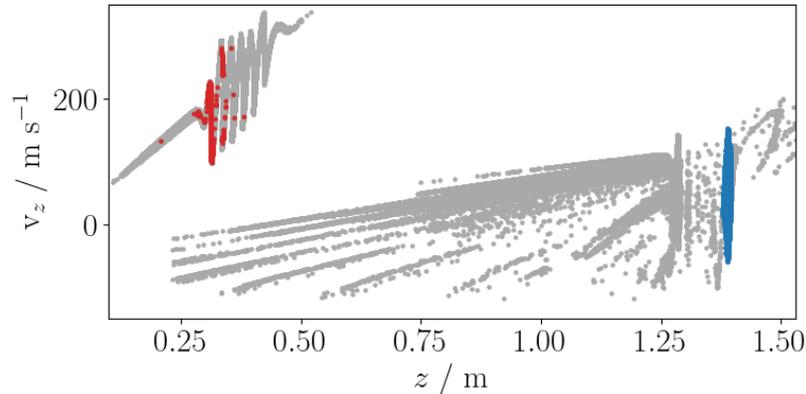


Figure 5.15: Deceleration of the  $F = 1, m_F = 1$  sub-level of  ${}^2S_{1/2}$  hydrogen, 9-module decelerator operating from 200 to 20  $\text{m s}^{-1}$  with peak current of 1000 A. Upper left, the atoms just as the second trap forms, lower / right, the atoms just as the trap collapses. The atoms located within the bounds of the trap at the end are in blue, their initial coordinates are in red.

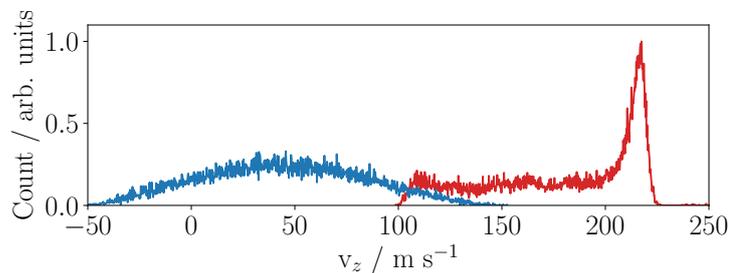


Figure 5.16: As per figure 5.15, but comparing the longitudinal velocities of the atoms remaining in trap 2 as it collapses (blue) with their velocities as the trap formed (red).

The trailing bunched atoms are marked in figure 5.17, and it can be seen that they originate from the lower-velocity tail of the beam. They were unintentionally loaded into a trap following the main target trap, but were subsequently still effectively bunched and decelerated. The number density of

the atoms in this cluster is lower, they only comprise  $\approx 0.2\%$  of the initial sample, and they have been decelerated from mean  $v_z = 134.0$  to  $57.8 \text{ m s}^{-1}$ . These plots suggest that the decelerator could effectively work with atomic hydrogen across a range of possible initial flow velocities and temperatures, and still produce a suitable output sample.

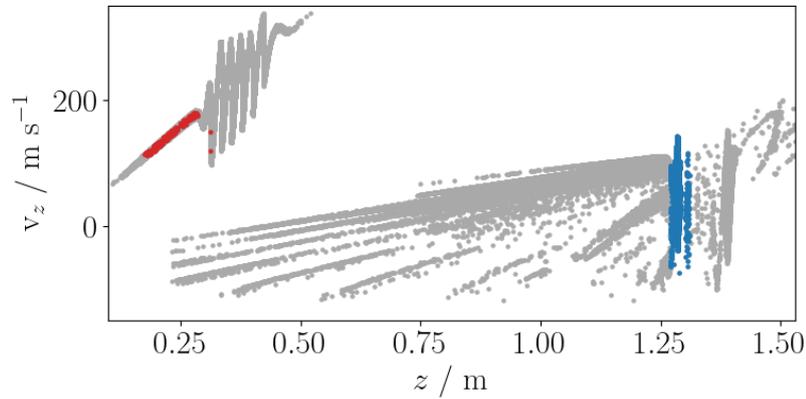


Figure 5.17: As per figure 5.15, but considering the other strongly-bunched cluster of atoms trailing those in trap 2.

One immediate point arising from consideration of figures 5.4 and 5.15 is that although the decelerated groups of particles are strongly spatially bunched, they have a large spread in velocity. This is a function of the deep traps arising from the large peak currents set in these simulations, so that the longitudinal separatrices are large. It may assist the loading of the MT-MOT to choose a smaller peak current and therefore get bunches of atoms with a reduced spread of velocity. The output of such a simulation is shown in figures 5.18 and 5.19. The bunch of atoms in trap 2 comprise  $\approx 4.6\%$  of the sample, and their mean velocity in  $z$  has decreased from  $179.9$  to  $36.1 \text{ m s}^{-1}$ .

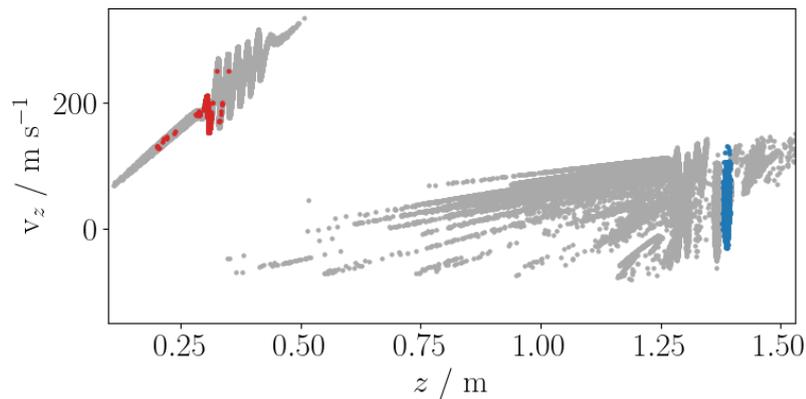


Figure 5.18: As per figure 5.15, but decelerator operating at peak current of 500 A.

The phase-space distribution of the atoms in the end traps for the two current values are compared in figure 5.20. This gives some interesting results. The simulation on the left has one bunch of atoms with a mean longitudinal velocity  $\bar{v}_z = 44.1 \text{ m s}^{-1}$  and standard deviation  $\sigma_{v_z} = 38.5 \text{ m s}^{-1}$ . The lower-current simulation on the right has two clusters; the rightmost, corresponding to the bunch

in the high-current simulation, has  $\bar{v}_z = 36.1 \text{ m s}^{-1}$  and  $\sigma_{v_z} = 26.4 \text{ m s}^{-1}$ . This is in line with our hypothesis that shallower decelerator traps should produce a packet of atoms with a smaller velocity spread, with the added bonus that it is also moving more slowly. What is even more significant is the density of the two packets; the high-current simulation produced a cluster comprising  $\approx 2.7 \%$  of the initial population of hydrogen atoms, whereas the equivalent bunch in the shallower-trap run contains  $\approx 4.6 \%$ . Although the trailing packet in the right plot has a similar shape and velocity distribution it is less dense, and only contains  $\approx 0.2 \%$  of the initial hydrogen.

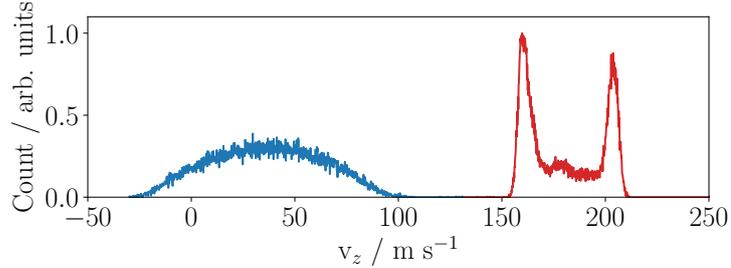


Figure 5.19: As per figure 5.18, but comparing the longitudinal velocities of the atoms remaining in trap 2 as it collapses (blue) with their velocities as the trap formed (red).

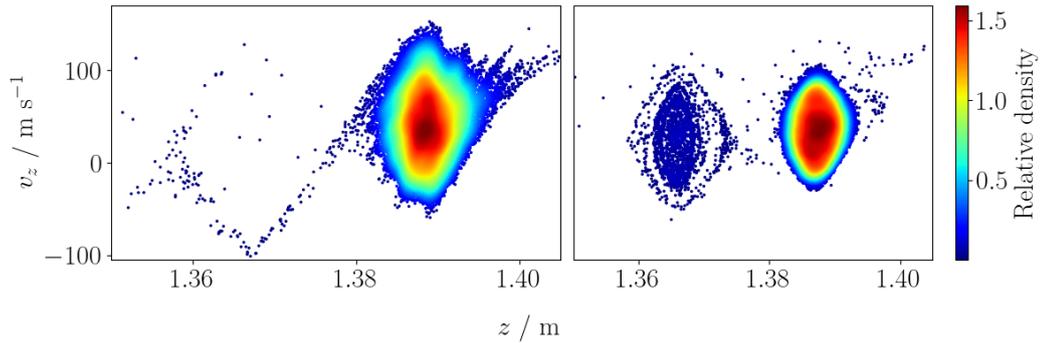


Figure 5.20: Hydrogen deceleration, shown are the atoms bunched in the end traps as the field collapses. Left, peak decelerator current of 1000 A, right, 500 A. The colour gradient shows the relative density of atoms.

A similar analysis for lithium is shown in figure 5.21; encouragingly, this also gave very positive results. The single main bunch of atoms in the high-current plot had  $\bar{v}_z = 44.5 \text{ m s}^{-1}$  and  $\sigma_{v_z} = 38.6 \text{ m s}^{-1}$ , and its equivalent in the low-current simulation had  $\bar{v}_z = 36.4 \text{ m s}^{-1}$  and  $\sigma_{v_z} = 26.4 \text{ m s}^{-1}$ . The relative proportion of the initial sample has increased from  $\approx 5.9 \%$  to  $\approx 9.7 \%$ . This indicates the importance of testing different combinations of experimental parameters, beyond simply adjusting the relative trigger times of the source and decelerator, and that simply setting the deceleration fields to their maximum value may not always be the optimal solution. If attempting to produce a dual-species beam comprising two species with a greater mass difference than the H-Li system considered here, it may be more difficult to find a set of parameters that provide a good compromise for deceleration and number density for both.

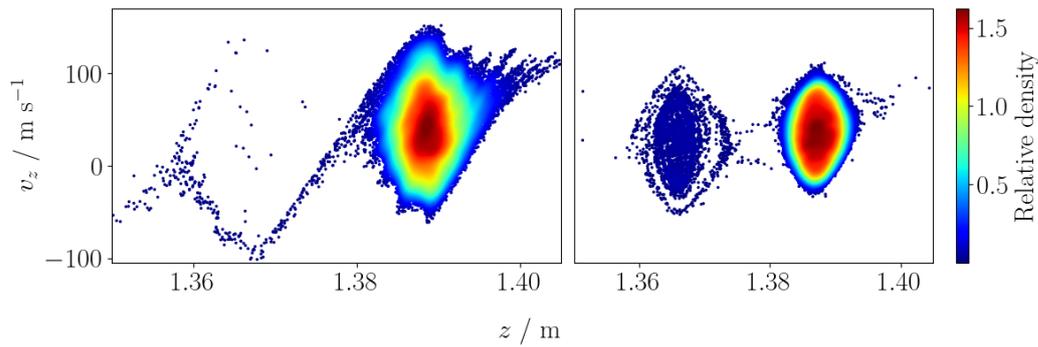


Figure 5.21: Lithium deceleration, shown are the atoms bunched in the end traps as the field collapses. Left, peak decelerator current of 1000 A, right, 500 A. The colour gradient shows the relative density of atoms. The distribution is very similar to that of hydrogen in figure 5.20, which is encouraging for the possibility of a decelerating a dual-species beam.

We had originally planned to see if we could create atomic hydrogen in our existing source.  $\text{H}_2$  has a dissociation energy of 4.52 eV so it seems possible that we can produce atomic hydrogen via the dielectric barrier discharge unit in our Even-Lavie valve; we currently produce metastable  $^3\text{P}_2$  argon, which has an energy of 11.54 eV. We constructed a circulation chamber in which  $\text{H}_2$  was to be mixed with xenon before being fed into the valve.

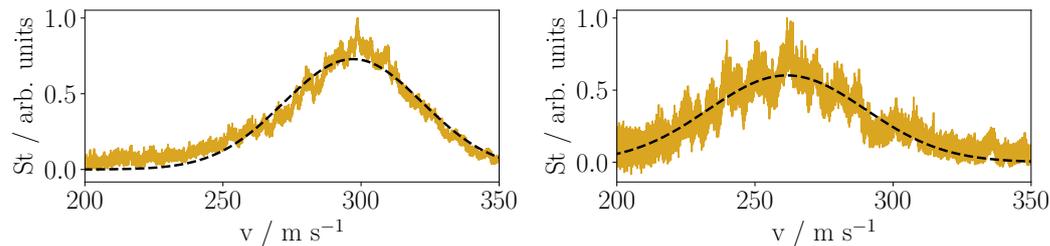


Figure 5.22: Supersonic expansion of xenon in the Even-Lavie valve with dielectric barrier discharge, detected on the microchannel-plate detector. This is converted to a speed distribution and fitted with equation 3.2.7, see section 3.2. Left / right, the valve is cooled to 200 / 180 K, giving fitted flow velocity of 290.6 / 251.9  $\text{m s}^{-1}$  and temperature of 10.1 / 13.3 K.

We began construction of a Wiley-McLaren mass spectrometer to be installed onto the end of the decelerator, interposed before the microchannel-plate detector. We intended to test the spectrometer by attempting to detect xenon; it should be possible to ionise metastable  $^3\text{P}_2$  xenon (excited in the DBD) using a YAG laser, either quadrupled to 266 nm or possibly the third harmonic at 355 nm. Alternatively, Gilmour *et al.* [243] describe detection of xenon in a Wiley-McLaren mass spectrometer via a two-photon dipole-forbidden transition driven at 249 nm, which should be achievable with a dye laser, which we also possess. Another possibility for producing atomic hydrogen is the scheme described in Hogan *et al.* [244]; ammonia is fed through a capillary into the region next to the pulsed valve where it is photolysed by an excimer laser operating at 193 nm. The resultant hydrogen atoms are then entrained into the supersonic expansion of the carrier gas from the valve. This is a proven technique but would require non-trivial reworking of our source chamber, to say nothing of the issues associated with handling ammonia. In either case, the hydrogen would be ionised for detection using the dye laser operating at 243 nm.

Unfortunately we did not get the opportunity to undertake this work. We did at least produce a beam of xenon with the Even-Lavie valve and detect it at the MCP. Figure 5.22 shows two speed distributions derived from time-of-flight signals, using the methods described in section 3.2, and fitted with equation 3.2.7 in order to extract estimates of the flow velocity  $v_z$  and temperature  $T_T$ . The source was operating with a stagnation pressure of 1.9 bar and cooled to between  $\approx 180$  K to 200 K, giving flow velocities of  $\approx 250$  m s $^{-1}$  to 290 m s $^{-1}$ . Only a small amount of data was taken and so the source parameters were not at all well optimised, so it seems likely that even slower beams should be achievable with some adjustments. It remains to be tested, however, how feasible it would be to entrain hydrogen into such a beam. There is a concern that the mixture of such light and heavy species would give rise to velocity slip.

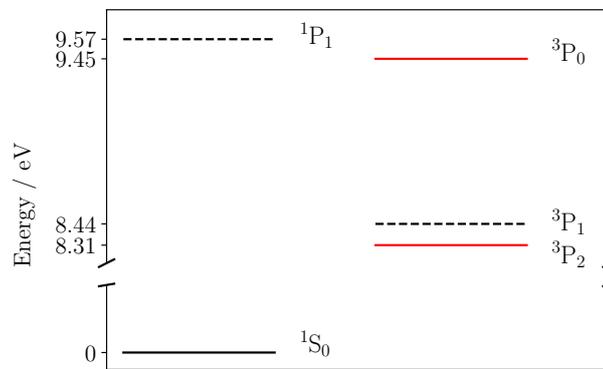


Figure 5.23: Energy levels of xenon in the first excited configuration, relative to the ground state. Dashed lines indicate states with a short excited lifetime before undergoing radiative decay, states in red are metastable.

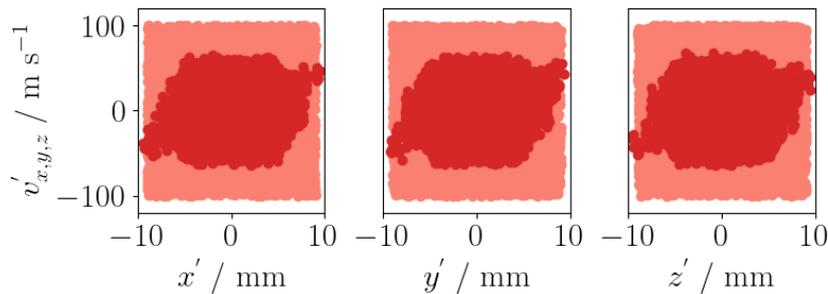


Figure 5.24: Acceptance of the  $F = 1, m_F = 1$  sub-level of  $^2S_{1/2}$  hydrogen in the MT-MOT. Pink, initial atom distribution, red, those remaining in the trap field after 2 ms.

We did not attempt to guide or decelerate the xenon beam. It seems possible that we could adjust the DBD parameters such that all metastable xenon atoms produced would be in the  $^3P_2$  state; unlike in argon, the two metastable states are separated by a large energy gap. The energies of the first four excited states of xenon are shown in figure 5.23. It seems unlikely that the 4-module decelerator could significantly operate on a species with such an unfavourable ratio of magnetic moment to mass, but it may be worth trying if it was desirable to attempt to quantify the relative state population of the

beam. Finally, the acceptance of the MT-MOT for hydrogen was assessed, as shown in figure 5.24, and was found to be  $1.1 \times 10^8$  (mm m s<sup>-1</sup>)<sup>3</sup>.

Other groups have previously decelerated dual-species beams, for example Akerman *et al.* describe doing so with oxygen and metastable argon [56], using their travelling-wave Zeeman decelerator (referred to here as the ‘Tel Aviv’ decelerator for convenience). In section 4.6 we considered the estimated acceptance of our decelerator for both argon and oxygen. The results for the two species were comparable, suggesting that the Durham decelerator could work with such a dual-species beam, but at the high beam velocities described by Akerman the acceptance was low in comparison to the Tel Aviv machine. Even with the proposed 9-module extension discussed here, we would not be able to successfully decelerate these species from such initial conditions.

Akerman *et al.* also report decelerating and then trapping a beam of oxygen and lithium [57]. Both species were loaded into a trap formed from two permanent magnets, producing a quadrupole trap geometry. This was later replaced with a superconducting trap, and used to capture lithium and oxygen [58] and then a co-decelerated beam of carbon atoms and molecular oxygen [59].

Our preliminary results with lithium are more encouraging, with an estimated acceptance that is 1 – 2 orders of magnitude larger than we found for oxygen or argon, albeit again we are considering a much slower initial beam flow velocity. The expected emittance of the decelerator for these 2 species seems to be comparable to the acceptance of the MT-MOT, which could be confirmed by further simulations using the methods described previously by the author [151] and in this thesis. Compared to our previous simulations, our quadcubic interpolator has improved our trap-loading models.

In conclusion; the proposed extended decelerator and MT-MOT show promise for decelerating a dual-species beam of Li and H, with the caveat that significant uncertainties surround the properties of the input to the decelerator. Further work should try to generate more plausible initial conditions for the two species to be loaded, and work to optimise and assess their deceleration and loading into the trap could be done.

The hydrogen and lithium decelerator acceptance estimates given for the decelerator are all with the wire quadrupole active. In the case of argon, sections 4.4 and 4.5.2 described the surprising result that the permanent magnet quadrupole did not always provide greater acceptance under all decelerator operations, and that is also the case for hydrogen and lithium. Although he did not see it for argon, McArd also described this result for hydrogen and lithium [2]. It is clear that further work needs to be done to investigate these effects and amend the design of the permanent magnet quadrupole, especially since its poorer performance appears to be worst for lighter species.

## 5.4 Buffer Gas Source

The Even-Lavie valve and dielectric-barrier discharge source have been described in sections 2.2.2 and 2.2.3. Despite their several advantages, a possible alternative is to replace them with a buffer gas source. This should offer comparable output translational temperatures of the order of  $\approx 5$  K but lower flow velocities in the region of 50 to 200 m s<sup>-1</sup> [12]. This would enable slower beams of various species from the Zeeman decelerator and facilitate the loading of the MT-MOT. Buffer gas sources are typically populated via capillary loading or laser ablation, which offer high particle number densities. However, compared to a source such as the Even-Lavie valve the temporal widths are large [245], which can necessitate a careful consideration of the geometry of the source.

Such a source could be modelled on that described in Bulleid [246], based around a helium cryocooler. The production of atomic hydrogen was briefly considered in section 5.3.3, the associated lithium might be loaded into the cell by laser ablation from a solid target. Design of such a source and its compatibility with the travelling-wave decelerator should be explored in future work.

## 5.5 Non-Adiabatic Losses

If a perturbation to a quantum system acts slowly enough, then according to the adiabatic theorem it remains in its instantaneous eigenstate [247]. If particles in a magnetic trap are not able to adiabatically follow the magnetic field, either due to their motion or changes in the field, then they may undergo a transition from a low-field-seeking state to a high-field seeking state, and be ejected from the trap. These are known as ‘Majorana transitions’ [62] or ‘spin-flips’. Utilising the Born-Oppenheimer approximation the internal and centre-of-mass coordinates  $\mathbf{r}$  of the particle can be separated [248], giving adiabatic energy states at the latter with eigenstates  $|\phi_k\rangle$  and eigenenergies  $E_k$ :

$$H(\mathbf{r}) |\phi_k(\mathbf{r})\rangle = E_k(\mathbf{r}) |\phi_k(\mathbf{r})\rangle \quad (5.5.1)$$

Then, transition from the trapped states  $|\phi_t\rangle$  to the untrapped states  $|\phi_u\rangle$  has a maximum probability of [247]:

$$P_{max} \leq \left[ \frac{\langle \phi_t | d/dt | \phi_u \rangle}{|E_t - E_u|/\hbar} \right]^2. \quad (5.5.2)$$

A condition for adiabatic following of the trapping field can then be given by:

$$|\langle \phi_t | \frac{d}{dt} | \phi_u \rangle| \ll \frac{|E_t - E_u|}{\hbar}. \quad (5.5.3)$$

An equivalent way to present this condition is as  $\omega_T \ll \omega_Z$ ; here  $\omega_Z$  is the Larmor precession rate, given by  $\omega_Z = \mu_B B/\hbar$ , and  $\omega_T$  is the rate at which the orientation of the field  $B$  changes [76]. In previous work tests of this condition were applied to the MT-MOT [236, 151] to estimate the likelihood of Majorana losses for lithium. At each timestep the code determines the magnetic field at the position of each trapped particle; as usual it uses the gradient of the field to calculate the magnetic trapping force, but the vector components of the field are also interpolated. The angle between this vector and the value that was found at the previous position of the particle is calculated, and this can

be used with the timestep to estimate the angular frequency of the particle motion. Likewise, the magnitude of the field is interpolated at each point, this can be averaged with the previous value to estimate the average Larmor frequency across the timestep.

The ratio of Larmor to orbital angular frequency can then be used as a rough guide to how likely a spin-flip transition would be. However, the question is; how large does the ratio have to be for us to assume that no transitions occur? In the earlier simulations of the MT-MOT we applied various threshold values of this ratio and ran simulations to count the percentage of particles lost for each. Although imprecise, this gave some suggestive information as to how bad any problems were likely to be, and the results suggested that the trap might not perform as poorly as had been feared [225].

A version of this code was adapted to test the decelerator field, using a similar method to that described for the acceptance tests in section 4.3; a ‘snapshot’ of the decelerator field around one single trap was extracted and populated with a uniform sample of atoms. This was then evolved under magnetic trapping, with the ratio of frequencies calculated as described above. Rather than remove particles, as in the earlier work, the mean of the ratio was saved. The simulation ran for a virtual time of 500  $\mu\text{s}$ .

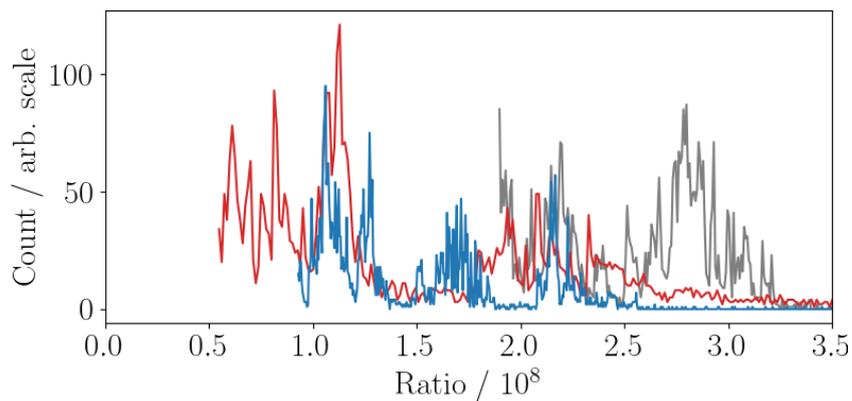


Figure 5.25: Estimated ratio of Larmor precession rate to orbital angular frequency in the decelerator field, guiding mode, peak current 500 A. Particles were evolved for simulated 500  $\mu\text{s}$ . Grey, argon, red, hydrogen, blue, lithium.

Figure 5.25 shows histograms of the mean ratio for argon (grey), hydrogen (red) and lithium (blue). Although this is only a very rough test, it seems that non-adiabatic losses should not be significant for this system, since the approximate adiabatic condition is generally met by at least a factor of  $10^6$  for all species. This does not appear to vary too much under deceleration, when the particles explore different regions of the trap; this is shown in figures 5.26 and 5.27.

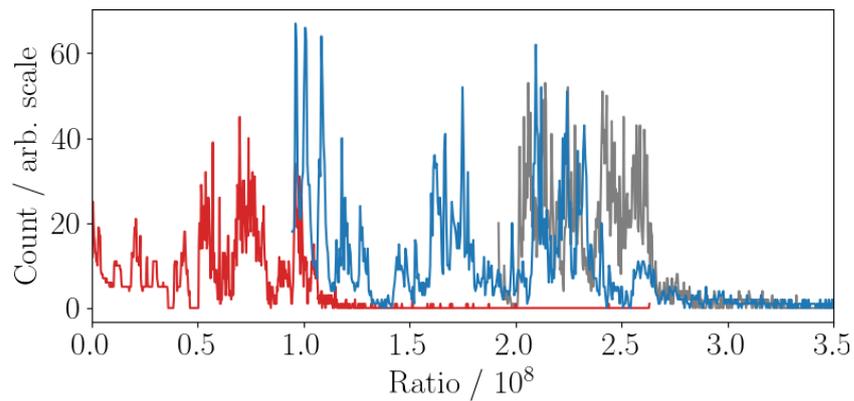


Figure 5.26: As per figure 5.25 but decelerating from 350 to 300 m s<sup>-1</sup>.

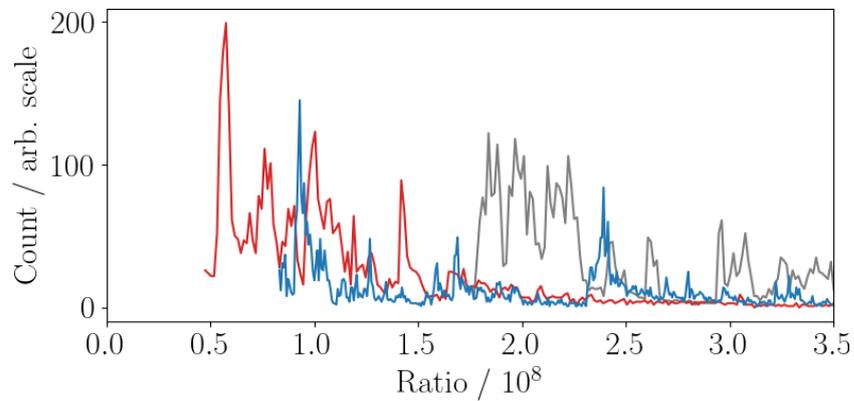


Figure 5.27: As per figure 5.25 but decelerating from 350 to 250 m s<sup>-1</sup>.

A version of the fully time-dependent decelerator code was developed to incorporate these tests, but not pursued given the preliminary results given by the static code. However, it is worth noting that Trimeche reported seeing a disagreement between the expected signal and data when decelerating metastable argon, using the decelerator upon which our experiment was based [222]. This was specifically the case for deceleration but not constant-velocity guiding, and the weakened signal was attributed to losses via non-adiabatic transitions. Possibly this is due to the different shape of the phase-stable region, seen in *e.g.* figure 4.30; the particles with higher velocities spend more time in the region of low field in the centre of the trap. A similar effect cannot be ruled out for our experiment without more analysis. Future work could attempt to calculate the actual transition probabilities for this system, and indeed the MT-MOT. For the latter, Majorana losses may be more significant due to the lower field magnitudes and the time the particles may spend in the region of the trap centre.

# Chapter 6

## Conclusion

This thesis has been largely concerned with the development of a comprehensive toolkit of software tools used to analyse and explain our travelling-wave Zeeman decelerator. This experiment produces data that can be difficult to understand without utilising Monte Carlo techniques to reproduce and thus quantify the signals that we receive. Chapter 1 gave a very brief overview of the motivation and background for the intended research goals for the project. Some consideration of the context and alternative approaches to working with cold molecules were given.

Chapter 2 gave a brief review of some of the technical aspects of the project that were previously undertaken [1, 2]. It also included more detailed studies of the specific aspects of the experiment that have been developed in this thesis, with one of the most significant being the work on supersonic expansions and time-of-flight data. This is explored further in chapter 3, which deals with the specific applications of computer modelling to the experiment. A careful analysis of the theoretical structure of gas expansions in a source with small temporal and spatial widths has been used to generate methods for simulating these expansions which are remarkably close to those we record in our experimental data. These methods might have some application for other groups using slightly different apparatus but that could be explored in future work.

Chapter 3 continued with descriptions of our methods for calculating the complex time-dependent magnetic fields used in the decelerator. Key to our ability to utilise these model fields in our simulations is our ability to rapidly and accurately interpolate them using the tricubic and quadcubic code libraries that we have developed. These are perhaps the most significant outcome of the work in this thesis, having already spawned several external collaborations. It is only in the relatively recent past that easy access to computers powerful enough to undertake large-scale numerical calculations has become very common, and these codes have arrived at an opportune time when other groups were looking for a solution of this type. Indeed, if suitable libraries had existed in 2018 then our interpolation codes would not have been written. Finally, the chapter concludes by comparing some argon data with simulations, allowing us to quantify the effects of the decelerator. The previously-unsuspected presence of other metastable terms of argon in the data is described and justified.

Chapter 4 takes a step back to review some of the ways in which the outputs of the developing simulation codes were fed back to make better estimates for the parameters that we should set when operating the experiment, specifically the combinations of settings for the power electronics and the source. Models for the optimum loading, by maximising the fraction of the incoming beam overlaid on the estimated longitudinal separatrix, are described. One of the most significant issues relating to this experiment is the effects of the PWM synthesis technique on the depth and position of the moving traps in the beam, and this is carefully analysed. Models for the acceptance of the decelerator are developed and justified, and compared with earlier models. A significant amount of work has gone into analysing and comparing the effects arising from the two possible quadrupole focussing fields,

but more work remains to be done. We have shown that the data given in chapter 3 was somewhat sub-optimal in the choice of experimental parameters used, and some predictions are made relating to argon. We have compared the capabilities of some other travelling-wave Zeeman decelerators with our experiment, by using our simulation codes to partially reproduce their published results.

Finally, chapter 5 gives some outlook for the project. It has always been known that the decelerator should be extended, and some simulations are made to predict the possibilities if this work were done (this was also considered in section 4.6). The proposed hybrid magnetic trap / magneto-optical trap is briefly described. The work undertaken by other members of the group on this project, utilising some of the resources developed in this thesis, shows promise. Dual-species beam simulations of hydrogen and lithium and their loading into the trap are encouraging but questions remain to be answered; specifically, how realistic are the estimated initial properties of the beam? These properties include the flow velocity, translational temperature, spatial distribution and number density of the relevant species. Given the importance of additional transverse focussing on a decelerator with our coil geometry, what is the best solution? The performance of the existing permanent magnet quadrupole design is not as clearly superior to the wire quadrupole as we had hoped. However, a wire quadrupole brings technical challenges, especially if the decelerator is lengthened.

The question of the decelerator length is also not straightforward to answer. Preliminary simulations indicate that simply extending the machine to match the size of comparable experiments would not allow us to match their existing results for comparable species and conditions. However, our design has potential for very strong confinement at lower beam velocities, which we have hardly explored.

It is unfortunate that further progress towards the research goals of the project has not been made. This is in no small part due to the 2020 pandemic which effectively stopped physical work in the laboratory for a prolonged time. The positive outcome to this is that the greater theoretical treatment of the project that this mandated has given us a vastly improved understanding of the characteristics and capabilities of the decelerator. We believe that we have demonstrated that the machine works as intended. Additionally, data obtained subsequent to the main work in this thesis are presented in Appendix A, and the analysis of these data is in good agreement with our expectations. This is despite the apparent presence of excessive unwanted  $m_J = 0$  argon atoms in the beam, which has been one of the most significant barriers to our characterisation of the decelerator. It would be desirable to address this in more detail, either by switching to a different species or different beam source. A longer decelerator would also help here, since non-trappable particles would be more likely to disperse.

Some discussion of the future prospects of the decelerator have been given. The current author and previous members of the group have done preliminary work on proposed experiments involving the co-deceleration and trapping of CaH with Li [2, 151] and H with Li [236, 241, 189]. Assuming that we were able to construct a suitable source, it seems entirely feasible that at least H and Li could be decelerated very well in our decelerator. We could also extend the existing simulation work on the MT-MOT to gain further understanding of the feasibility of sympathetic cooling of H or CaH, or indeed other species.

# Appendix A

## Additional Argon Data

Subsequent to the completion of this thesis there was an opportunity to perform an additional set of argon guiding and deceleration experiments with the Travelling-Wave Zeeman Decelerator, allowing us to test the predictive power of the simulation methods described in the preceding chapters. The results, given here, were somewhat better than the data presented in the main thesis.

Prior to capturing new data various combinations of experimental parameters were assessed via simulation. We had a good understanding of the expected flow velocity and translational temperature of the source, from analysis of our previous data. For a given initial travelling-wave velocity and peak current we were therefore able to estimate the electronics delay time that should lead to the greatest capture of atoms within the phase-stable region of the target trap. The decelerator wave velocity can be chosen to match the central portion of the beam, where it is most dense, or part of the slower tail.

Transmission (*i.e.* zero field) data were captured and fitted as previously described; the beam flow velocity was estimated to be in the region of  $380 - 390 \text{ m s}^{-1}$ . An initial wave velocity of  $380 \text{ m s}^{-1}$  was chosen, with final velocities of 380, 360 and  $340 \text{ m s}^{-1}$ . The decelerator power electronics were set with a peak current of 400 A at 600 V. For this combination of initial wave velocity and current the estimated optimal power electronics delay was  $725 \mu\text{s}$ ; additional data were taken in which the delay was swept by  $\pm 10 \mu\text{s}$  in  $2 \mu\text{s}$  increments.

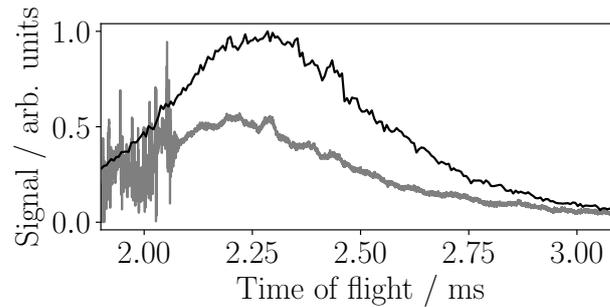


Figure A.1: Guiding of the  $m_J = 2$  sub-level of  $^3\text{P}_2$  argon, decelerator operating at  $380 \text{ m s}^{-1}$  with peak current of 400 A. Experimental data is in grey with simulated time-of-flight overlaid in black, the latter padded with additional  $^3\text{P}_0$ .

The Even-Lavie valve was configured with a  $25.8 \mu\text{s}$  pulse duration, and the argon feed was at 9 bar. The valve assembly was cooled to 138 K. The DBD was set to trigger  $165 \mu\text{s}$  after the valve, with RF voltage of 1200 V, RF current of 10 A, and 27 pulses at 700 kHz. Figure A.1 shows attempted constant velocity guiding at  $380 \text{ m s}^{-1}$ ; the experimental data are in grey with the simulated results overlaid in black. The simulation comprised  $1 \times 10^6$  argon atoms equally distributed amongst the 5 sub-levels of the  $^3\text{P}_2$  term; an additional  $8 \times 10^6$   $^3\text{P}_0$  were added to reproduce the shape of the data.

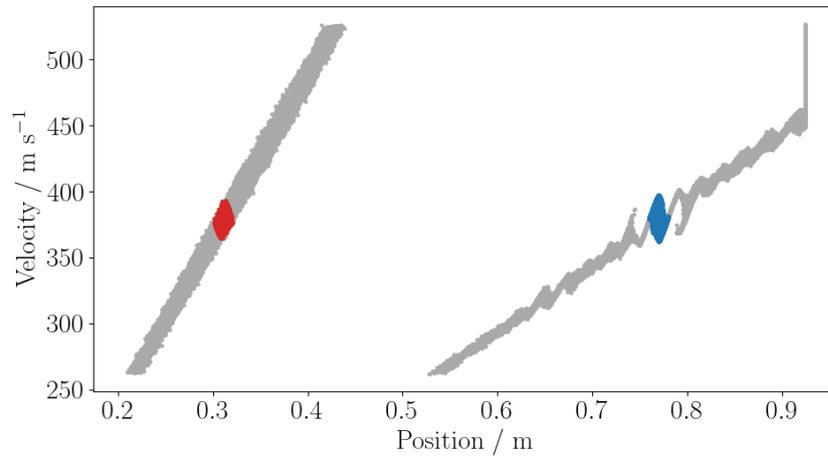


Figure A.2: Guiding of argon, trap velocity of  $380 \text{ m s}^{-1}$ , peak current of 400 A. Left, the atoms just as the second trap forms, right, the atoms just as the trap collapses. The atoms located within the bounds of the trap at the end are in blue, their initial coordinates are in red.

Figure A.2 shows the longitudinal phase-space coordinates of the simulated beam at the beginning and end of the guiding sequence - only the atoms in the  $m_J = 2$  state are shown. The atoms located between the maxima of the second trap of the decelerator wave as the field shuts down are marked in blue on the right, these same atoms are marked red in the left plot. The mean velocities of these atoms, plotted in figure A.3, increased from 378 to  $380 \text{ m s}^{-1}$ .

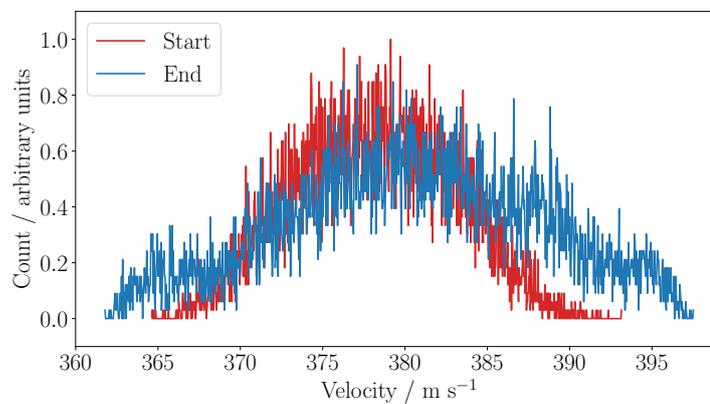


Figure A.3: As per figure A.2, but comparing the longitudinal velocities of the atoms remaining in trap 2 as it collapses (blue) with their velocities as the trap formed (red).

The broad, almost multi-peaked structure of the blue velocity plot in figure A.3 can be explained by figure A.4, which is a zoomed-in view of figure A.2. The separatrix is not completely filled, since for this short decelerator there has not been enough time for the atoms to rotate in phase-space. It is, however, more uniformly populated than the trap seen in *e.g.* figure 3.33 in section 3.10.

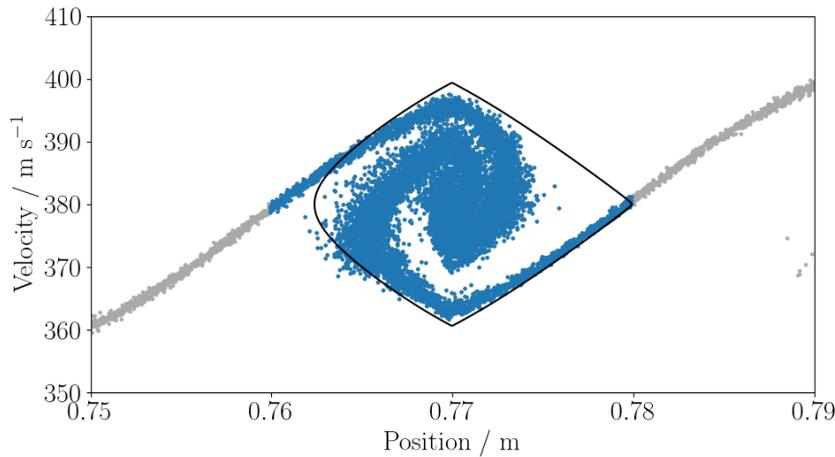


Figure A.4: As per figure A.2, but zoomed in on the atoms in trap 2 at the end of the guiding sequence. The estimated longitudinal separatrix is in black.

At the time the data were taken measurements of the current through several phases of the decelerator coils were also captured. Using the PWM technique we would expect the peak current to reach the target of 400 A to within  $\approx 10\%$ , however later analysis of the probe results showed that the peak current was instead reduced by up to 20%, although this was not consistent across all of the phases measured. At the time of writing we are not certain where the fault lies. Each data set was simulated at both 300 and 400 A, with the former value generally giving the best match between the data and prediction. The simulation code includes measured values for the resistance and inductance of each phase of each coil, ideally these should be physically tested again. The components of the power electronics should also be tested for faults, but it has not yet been possible to do so. Nonetheless, we see reasonably good agreement between the simulations and data presented in this section.

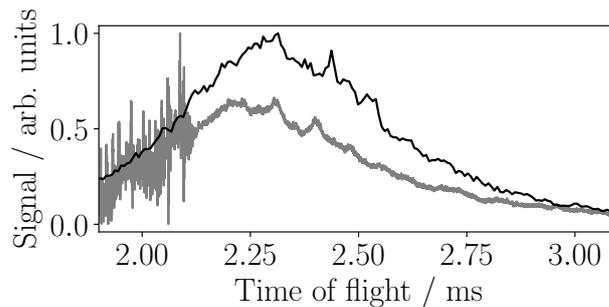


Figure A.5: Deceleration of the  $m_J = 2$  sub-level of  $^3P_2$  argon, decelerator operating from 379 to 363  $\text{m s}^{-1}$  with peak current of 400 A. Experimental data is in grey with simulated time-of-flight overlaid in black, the latter padded with additional  $^3P_0$ .

Deceleration from 380 to 360  $\text{m s}^{-1}$  was then attempted; in figure A.5 the experimental data are in grey with the simulated results overlaid in black. The simulation comprised  $1 \times 10^6$  argon atoms equally distributed amongst the 5 sub-levels of the  $^3P_2$  term; an additional  $4 \times 10^6$   $^3P_0$  were added to reproduce the shape of the data. For this example the power electronics delay was set to 731  $\mu\text{s}$ , and simulation of the decelerator wave shows that it achieved an initial / final velocity of 379 / 363  $\text{m s}^{-1}$ .

Figure A.6 shows the longitudinal phase-space coordinates of the simulated beam before and after deceleration; again, only the atoms in the  $m_J = 2$  state are shown. The longitudinal velocities of these atoms are histogrammed in figure A.7; the mean values decrease from 378 to 366  $\text{m s}^{-1}$ , however it can be seen that a higher-velocity population of non-phase-stable atoms are present, distorting the results.

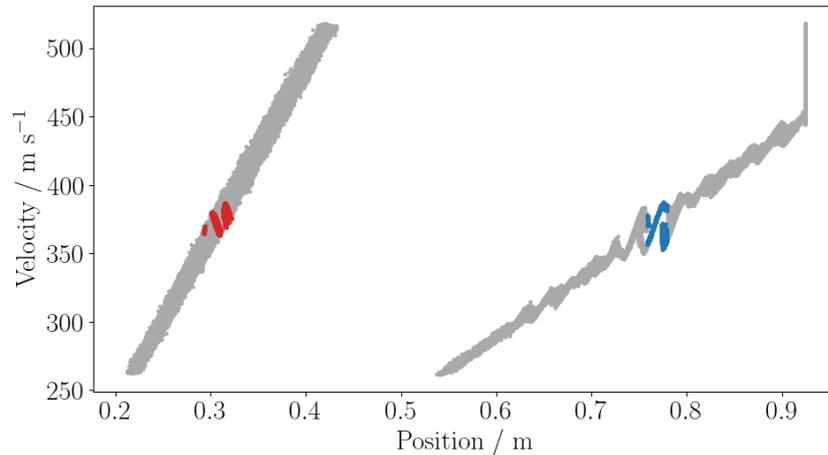


Figure A.6: Deceleration of argon, trap initial / final velocity of 379 / 363  $\text{m s}^{-1}$ , peak current of 400 A. Left, the atoms just as the second trap forms, right, the atoms just as the trap collapses. The atoms located within the bounds of the trap at the end are in blue, their initial coordinates are in red.

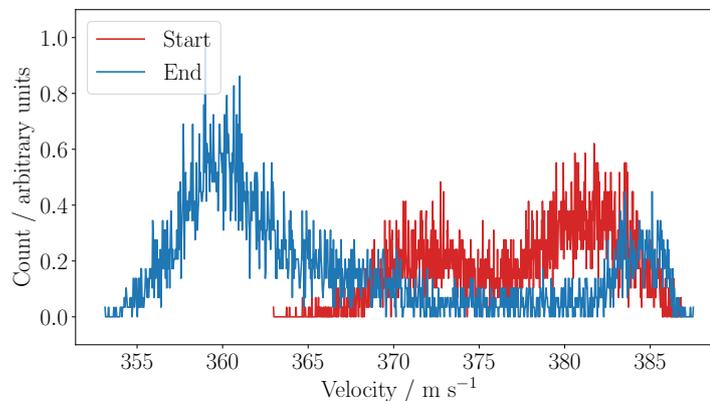


Figure A.7: As per figure A.6, but comparing the longitudinal velocities of the atoms remaining in trap 2 as it collapses (blue) with their velocities as the trap formed (red).

There is good agreement in the structure of the peaks between approximately  $2.3 \mu\text{s}$  to  $2.5 \mu\text{s}$ , with the specific peak relating to the  $m_J = 2$  atoms located in trap 2 being at  $\approx 2.48 \mu\text{s}$ . There is, however, disagreement in the central times of the peaks. Adjusting the simulated peak current shifts the relative positions of the peaks in time, but a value for which all features align with the data was not found. This supports our suspicion that the current drop was not constant across all coil phases.

Deceleration from  $380$  to  $350 \text{ m s}^{-1}$  was attempted, as shown in figures A.8 and A.9. The electronics delay was  $721 \mu\text{s}$  and  ${}^3\text{P}_0$  were added to the simulation in the ratio 5:1 with  ${}^3\text{P}_2$ . The time-of-flight is reproduced reasonably well, but as can be seen from the simulated phase-space plots the acceptance for this magnitude of deceleration is small, so that there are few atoms in the decelerated packet.

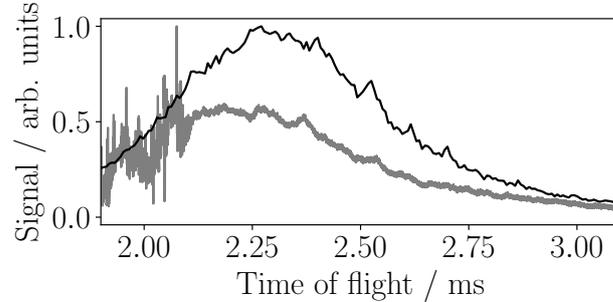


Figure A.8: Deceleration of the  $m_J = 2$  sub-level of  ${}^3\text{P}_2$  argon, decelerator operating from  $379$  to  $352 \text{ m s}^{-1}$  with peak current of  $400 \text{ A}$ . Experimental data is in grey with simulated time-of-flight overlaid in black, the latter padded with additional  ${}^3\text{P}_0$ .

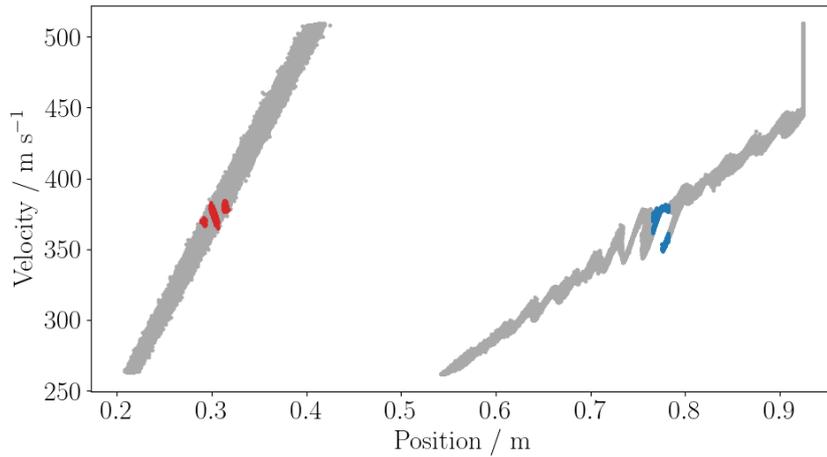


Figure A.9: Deceleration of argon, trap initial / final velocity of  $379 / 352 \text{ m s}^{-1}$ , peak current of  $400 \text{ A}$ . Left, the atoms just as the second trap forms, right, the atoms just as the trap collapses. The atoms located within the bounds of the trap at the end are in blue, their initial coordinates are in red.

An initial wave velocity of  $360 \text{ m s}^{-1}$  was then selected, with final velocities of  $360$ ,  $340$  and  $320 \text{ m s}^{-1}$ . Figure A.10 shows the experimental signal compared to simulated time-of-flight for deceleration from  $360$  to  $340 \text{ m s}^{-1}$  (simulation showed that the PWM technique achieved  $360$  to  $341 \text{ m s}^{-1}$ ). The power electronics delay was  $771 \mu\text{s}$ . In this case an additional  $5 \times 10^6$   ${}^3\text{P}_0$  atoms were added to the  $1 \times 10^6$  in the  ${}^3\text{P}_2$  term in order to match the bulk shape of the data. Although the peaks are weak in both the data and the simulation, there is good agreement in the region of interest, approximately between  $2.5$  and  $2.6 \text{ ms}$ . The longitudinal phase-space coordinates and velocities of the  $m_J = 2$  atoms are shown in figures A.11 and A.12, specifically those loaded into the second trap. Again, there is a presence of non-phase-stable atoms distorting the mean velocities, which are  $\approx 360 / 345 \text{ m s}^{-1}$  pre- / post-deceleration. The atoms in the phase-stable region can be identified, however, and show clear indication of deceleration by the intended amount.

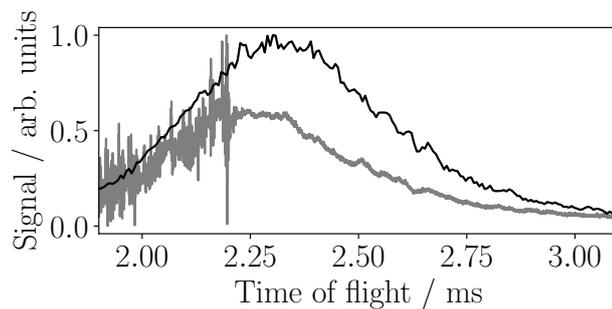


Figure A.10: Deceleration of the  $m_J = 2$  sub-level of  $^3P_2$  argon, decelerator operating from  $360$  to  $341 \text{ m s}^{-1}$  with peak current of  $400 \text{ A}$ . Experimental data is in grey with simulated time-of-flight overlaid in black, the latter padded with additional  $^3P_0$ .

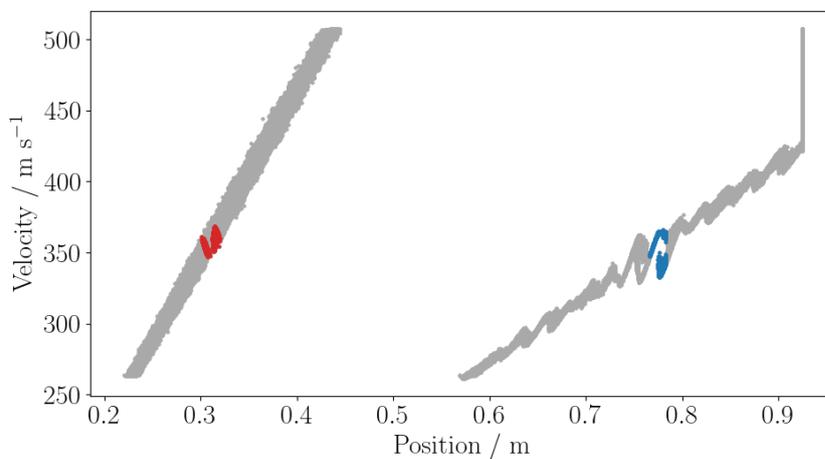


Figure A.11: Deceleration of argon, trap initial / final velocity of  $360 / 341 \text{ m s}^{-1}$ , peak current of  $400 \text{ A}$ . Left, the atoms just as the second trap forms, right, the atoms just as the trap collapses. The atoms located within the bounds of the trap at the end are in blue, their initial coordinates are in red.

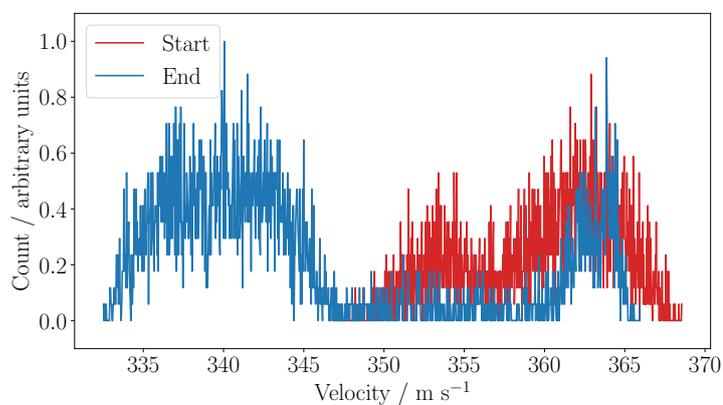


Figure A.12: As per figure A.11, but comparing the longitudinal velocities of the atoms remaining in trap 2 as it collapses (blue) with their velocities as the trap formed (red).

## Bibliography

- [1] A. Mizouri. 'A Moving-Trap Zeeman Decelerator'. PhD thesis. University of Durham, 2016.
- [2] L. McArd. 'A Travelling Wave Zeeman Decelerator For Atoms and Molecules'. PhD thesis. University of Durham, 2017.
- [3] K. Madison et al. *Annual Review of Cold Atoms and Molecules*. WORLD SCIENTIFIC, 2014. DOI: 10.1142/9100.
- [4] B. R. Heazlewood and T. P. Softley. 'Towards chemistry at absolute zero'. In: *Nature Reviews Chemistry* 5.2 (2021), pp. 125–140.
- [5] E. Altman et al. 'Quantum Simulators: Architectures and Opportunities'. In: *PRX Quantum* 2.1 (2021), p. 017003. DOI: 10.1103/PRXQuantum.2.017003.
- [6] M. T. Bell and T. P. Softley. 'Ultracold molecules and ultracold chemistry'. In: *Molecular Physics* 107.2 (2009), pp. 99–132. DOI: 10.1080/00268970902724955.
- [7] M. P. A. Jones, R. M. Potvliege and M. Spannowsky. 'Probing new physics using Rydberg states of atomic hydrogen'. In: *Phys. Rev. Res.* 2 (1 2020), p. 013244. DOI: 10.1103/PhysRevResearch.2.013244.
- [8] L. D. Carr et al. 'Cold and ultracold molecules: science, technology and applications'. In: *New Journal of Physics* 11.5 (2009), p. 055049.
- [9] T. P. Softley. 'Cold and ultracold molecules in the twenties'. In: *Proceedings of the Royal Society A* 479.2274 (2023), p. 20220806.
- [10] W. Demtröder. *Atoms, molecules and photons: an introduction to atomic-, molecular-, and quantum-physics*. Springer, 2011.
- [11] R. deCarvalho et al. 'Buffer-gas loaded magnetic traps for atoms and molecules: A primer'. In: *The European Physical Journal D-Atomic, Molecular, Optical and Plasma Physics* 7 (1999), pp. 289–309.
- [12] N. R. Hutzler, H.-I. Lu and J. M. Doyle. 'The buffer gas beam: An intense, cold, and slow source for atoms and molecules'. In: *Chemical reviews* 112.9 (2012), pp. 4803–4827.
- [13] S. Truppe et al. 'A buffer gas beam source for short, intense and slow molecular pulses'. In: *Journal of Modern Optics* 65.5-6 (2018), pp. 648–656. DOI: 10.1080/09500340.2017.1384516.
- [14] K. S. Twyman et al. 'Production of cold beams of ND<sub>3</sub> with variable rotational state distributions by electrostatic extraction of He and Ne buffer-gas-cooled beams'. In: *The Journal of Chemical Physics* 141.2 (July 2014), p. 024308. ISSN: 0021-9606. DOI: 10.1063/1.4885855.
- [15] J. Toscano et al. 'A magnetic guide to purify radical beams'. In: *The Journal of Chemical Physics* 149.17 (Nov. 2018), p. 174201. ISSN: 0021-9606. DOI: 10.1063/1.5053656.
- [16] S. C. Wright et al. 'Cryogenic buffer gas beams of AlF, CaF, MgF, YbF, Al, Ca, Yb and NO—a comparison'. In: *Molecular Physics* 121.17-18 (2023), e2146541.

- [17] Y. Takahashi et al. ‘Simulation of cryogenic buffer gas beams’. In: *Phys. Rev. Res.* 3 (2 2021), p. 023018. DOI: 10.1103/PhysRevResearch.3.023018.
- [18] R. Campargue. *Atomic and Molecular Beams*. Springer-Verlag, 2001.
- [19] K. N. C. Bray. ‘Elements of Gasdynamics. H. W. Liepmann and A. Roshko. Chapman and Hall, London, 1957. 439 pp. Illustrated.’ In: *The Journal of the Royal Aeronautical Society* 61.562 (1957), pp. 702–703. DOI: 10.1017/S0368393100132067.
- [20] C. Amarasinghe and A. G. Suits. ‘Intrabeam scattering for ultracold collisions’. In: *The journal of physical chemistry letters* 8.20 (2017), pp. 5153–5159.
- [21] W. E. Perreault, N. Mukherjee and R. N. Zare. ‘HD ( $v = 1, j = 2, m$ ) orientation controls HD–He rotationally inelastic scattering near 1 K’. In: *The Journal of Chemical Physics* 150.17 (May 2019), p. 174301. DOI: 10.1063/1.5096531.
- [22] R. Phaneuf et al. ‘Merged-beams experiments in atomic and molecular physics’. In: *Reports on Progress in Physics* 62.7 (1999), p. 1143.
- [23] Q. Wei, I. Lyuksyutov and D. Herschbach. ‘Merged-beams for slow molecular collision experiments’. In: *The Journal of chemical physics* 137.5 (2012).
- [24] A. B. Henson et al. ‘Observation of resonances in Penning ionization reactions at sub-kelvin temperatures in merged beams’. In: *Science* 338.6104 (2012), pp. 234–238.
- [25] M. Gupta and D. Herschbach. ‘A mechanical means to produce intense beams of slow molecules’. In: *The Journal of Physical Chemistry A* 103.50 (1999), pp. 10670–10673.
- [26] M. Gupta and D. Herschbach. ‘Slowing and speeding molecular beams by means of a rapidly rotating source’. In: *The Journal of Physical Chemistry A* 105.9 (2001), pp. 1626–1637.
- [27] S. Chervenkov et al. ‘Continuous centrifuge decelerator for polar molecules’. In: *Physical review letters* 112.1 (2014), p. 013001.
- [28] W. Gerlach and O. Stern. ‘Der experimentelle Nachweis der Richtungsquantelung im Magnetfeld’. In: *Zeitschrift für Physik* 9.1 (1922), pp. 349–352. ISSN: 0044-3328. DOI: 10.1007/BF01326983.
- [29] H. Bennewitz, W. Paul and C. Schlier. ‘Fokussierung polarer moleküle’. In: *Zeitschrift für Physik* 141 (1955), pp. 6–15.
- [30] S. Y. Van De Meerakker, H. L. Bethlem and G. Meijer. ‘Taming molecular beams’. In: *Nature Physics* 4.8 (2008), pp. 595–602.
- [31] H. L. Bethlem, G. Berden and G. Meijer. ‘Decelerating Neutral Dipolar Molecules’. In: *Phys. Rev. Lett.* 83 (8 1999), pp. 1558–1561. DOI: 10.1103/PhysRevLett.83.1558.
- [32] H. L. Bethlem et al. ‘Trapping Neutral Molecules in a Traveling Potential Well’. In: *Phys. Rev. Lett.* 84 (25 2000), pp. 5744–5747. DOI: 10.1103/PhysRevLett.84.5744.
- [33] H. L. Bethlem et al. ‘Electrostatic trapping of ammonia molecules’. In: *Nature* 406.6795 (2000), pp. 491–494.
- [34] L. Scharfenberg et al. ‘Operation of a Stark decelerator with optimum acceptance’. In: *Phys. Rev. A* 79 (2 2009), p. 023410. DOI: 10.1103/PhysRevA.79.023410.
- [35] S. A. Meek et al. ‘Trapping Molecules on a Chip in Traveling Potential Wells’. In: *Phys. Rev. Lett.* 100 (15 2008), p. 153003. DOI: 10.1103/PhysRevLett.100.153003.

- [36] A. Osterwalder et al. ‘Deceleration of neutral molecules in macroscopic traveling traps’. In: *Physical Review A* 81.5 (2010), pp. 1–4. ISSN: 10502947. DOI: 10.1103/PhysRevA.81.051401.
- [37] N. Bulleid et al. ‘Traveling-wave deceleration of heavy polar molecules in low-field-seeking states’. In: *Physical Review A* 86.2 (2012), p. 021404.
- [38] J. Van den Berg et al. ‘Traveling-wave deceleration of SrF molecules’. In: *Journal of Molecular Spectroscopy* 300 (2014), pp. 22–25.
- [39] M. Quintero-Pérez et al. ‘Static trapping of polar molecules in a traveling wave decelerator’. In: *Physical Review Letters* 110.13 (2013), p. 133003.
- [40] Y. Shyur, J. A. Bossert and H. J. Lewandowski. ‘Pulsed operation of a ring Stark decelerator’. In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 51.16 (2018), p. 165101. DOI: 10.1088/1361-6455/aad1b0.
- [41] J. Greenberg et al. ‘Velocity-tunable beam of continuously decelerated polar molecules for cold ion-molecule reaction studies’. In: *Review of Scientific Instruments* 92.10 (Oct. 2021), p. 103202. ISSN: 0034-6748. DOI: 10.1063/5.0057859.
- [42] S. D. Hogan, M. Motsch and F. Merkt. ‘Deceleration of supersonic beams using inhomogeneous electric and magnetic fields’. In: *Physical Chemistry Chemical Physics* 13.42 (2011), pp. 18705–18723.
- [43] G. Meijer. *Molecular beams in physics and chemistry: from Otto Stern’s pioneering exploits to present-day feats*. Ed. by B. Friedrich and H. Schmidt-Böcking. Springer Nature, 2021, p. 463.
- [44] C. V. Heer. ‘Feasibility of Containment of Quantum Magnetic Dipoles’. In: *Review of Scientific Instruments* 34.5 (1963), pp. 532–537. DOI: 10.1063/1.1718426.
- [45] K.-J. Kügler, W. Paul and U. Trinks. ‘A magnetic storage ring for neutrons’. In: *Physics Letters B* 72.3 (1978), pp. 422–424.
- [46] A. L. Migdall et al. ‘First Observation of Magnetically Trapped Neutral Atoms’. In: *Phys. Rev. Lett.* 54 (24 1985), pp. 2596–2599. DOI: 10.1103/PhysRevLett.54.2596.
- [47] N. Vanhaecke et al. ‘Multistage Zeeman deceleration of hydrogen atoms’. In: *Phys. Rev. A* 75 (3 2007), p. 031402. DOI: 10.1103/PhysRevA.75.031402.
- [48] E. Narevicius et al. ‘An atomic coilgun: Using pulsed magnetic fields to slow a supersonic beam’. In: *New Journal of Physics* 9.10 (2007), p. 358. ISSN: 13672630. DOI: 10.1088/1367-2630/9/10/358.
- [49] K. Dulitz. ‘Towards the study of cold chemical reactions using Zeeman decelerated supersonic beams’. DPhil thesis. Oxford University, 2014.
- [50] T. Cremers et al. ‘Multistage Zeeman decelerator for molecular-scattering studies’. In: *Phys. Rev. A* 95 (4 2017), p. 043415. DOI: 10.1103/PhysRevA.95.043415.
- [51] A. Trimeche et al. ‘Trapping of a supersonic beam in a traveling magnetic wave’. In: *The European Physical Journal D* 65.1 (2011), pp. 263–271. ISSN: 1434-6079. DOI: 10.1140/epjd/e2011-20096-1.
- [52] E. Lavert-Ofir et al. ‘A moving magnetic trap decelerator: A new source of cold atoms and molecules’. In: *New Journal of Physics* 13.103030 (2011). ISSN: 13672630. DOI: 10.1088/1367-2630/13/10/103030.

- [53] T. Damjanović et al. ‘A new design for a traveling-wave Zeeman decelerator: I. Theory’. In: *New Journal of Physics* 23.10 (2021), p. 105006. DOI: 10.1088/1367-2630/ac2b52.
- [54] Y. Liu and L. Luo. ‘Simultaneous Zeeman deceleration of polyatomic free radical with lithium atoms’. In: *Frontiers of Physics* 16 (2021), pp. 1–7.
- [55] Y. Ji et al. ‘Design of a ring-shaped traveling-wave Zeeman decelerator for both light and heavy molecules’. In: *Phys. Rev. A* 108 (4 2023), p. 043115. DOI: 10.1103/PhysRevA.108.043115.
- [56] N. Akerman et al. ‘Simultaneous deceleration of atoms and molecules in a supersonic beam’. In: *New Journal of Physics* 17.6 (2015). ISSN: 13672630. DOI: 10.1088/1367-2630/17/6/065015.
- [57] N. Akerman et al. ‘Trapping of Molecular Oxygen together with Lithium Atoms’. In: *Phys. Rev. Lett.* 119 (7 2017), p. 073204. DOI: 10.1103/PhysRevLett.119.073204.
- [58] Y. Segev et al. ‘Collisions between cold molecules in a superconducting magnetic trap’. In: *Nature* 572.7768 (2019), pp. 189–193.
- [59] M. Karpov et al. ‘Low-energy collisions between carbon atoms and oxygen molecules in a magnetic trap’. In: *New Journal of Physics* 22.10 (2020), p. 103055. DOI: 10.1088/1367-2630/abc391.
- [60] P. Jansen and F. Merkt. ‘Manipulating beams of paramagnetic atoms and molecules using inhomogeneous magnetic fields’. In: *Progress in Nuclear Magnetic Resonance Spectroscopy* 120-121 (2020), pp. 118–148. ISSN: 0079-6565. DOI: 10.1016/j.pnmrs.2020.08.002.
- [61] P. Jansen, L. Semeria and F. Merkt. ‘High-resolution spectroscopy of He2+ using Rydberg-series extrapolation and Zeeman-decelerated supersonic beams of metastable He2’. In: *Journal of Molecular Spectroscopy* 322 (2016), pp. 9–17. ISSN: 0022-2852. DOI: 10.1016/j.jms.2016.01.013.
- [62] E. Majorana. ‘Atomi orientati in campo magnetico variabile’. In: *Nuovo Cimento* 9 (1932), pp. 43–50.
- [63] M. H. Anderson et al. ‘Observation of Bose-Einstein condensation in a dilute atomic vapor’. In: *science* 269.5221 (1995), pp. 198–201.
- [64] K. B. Davis et al. ‘Bose-Einstein condensation in a gas of sodium atoms’. In: *Physical review letters* 75.22 (1995), p. 3969.
- [65] B. K. Stuhl et al. ‘Evaporative cooling of the dipolar hydroxyl radical’. In: *Nature* 492 (2012).
- [66] G. Valtolina et al. ‘Dipolar evaporation of reactive molecules to below the Fermi temperature’. In: *Nature* 588.7837 (2020), pp. 239–243.
- [67] J. Schmidt et al. ‘Optical traps for sympathetic cooling of ions with ultracold neutral atoms’. In: *Physical review letters* 124.5 (2020), p. 053402.
- [68] T. Baba and I. Waki. ‘Sympathetic cooling rate of gas-phase ions in a radio-frequency-quadrupole ion trap’. In: *Applied Physics B* 74 (2002), pp. 375–382.
- [69] T. V. Tscherbul, H.-G. Yu and A. Dalgarno. ‘Sympathetic Cooling of Polyatomic Molecules with S-State Atoms in a Magnetic Trap’. In: *Phys. Rev. Lett.* 106 (7 2011), p. 073201. DOI: 10.1103/PhysRevLett.106.073201.

- [70] J. Lim et al. ‘Modeling sympathetic cooling of molecules by ultracold atoms’. In: *Phys. Rev. A* 92 (5 2015), p. 053419. DOI: 10.1103/PhysRevA.92.053419.
- [71] M. Warehime and J. Kłos. ‘Nonadiabatic collisions of CaH with Li: Importance of spin-orbit-induced spin relaxation in spin-polarized sympathetic cooling of CaH’. In: *Phys. Rev. A* 92 (3 2015), p. 032703. DOI: 10.1103/PhysRevA.92.032703.
- [72] H. Son et al. ‘Collisional cooling of ultracold molecules’. In: *Nature* 580.7802 (2020), pp. 197–200.
- [73] M. D. Frye et al. ‘Approach to chaos in ultracold atomic and molecular physics: Statistics of near-threshold bound states for Li+CaH and Li+CaF’. In: *Phys. Rev. A* 93 (5 2016), p. 052713. DOI: 10.1103/PhysRevA.93.052713.
- [74] M. Morita, R. V. Krems and T. V. Tscherbul. ‘Universal Probability Distributions of Scattering Observables in Ultracold Molecular Collisions’. In: *Phys. Rev. Lett.* 123 (1 2019), p. 013401. DOI: 10.1103/PhysRevLett.123.013401.
- [75] W. D. Phillips. ‘Nobel Lecture: Laser cooling and trapping of neutral atoms’. In: *Reviews of Modern Physics* 70.3 (1998), p. 721.
- [76] H. Metcalf and P. van der Straten. *Laser Cooling and Trapping*. Springer-Verlag, 1999.
- [77] E. S. Shuman, J. F. Barry and D. DeMille. ‘Laser cooling of a diatomic molecule’. In: *Nature* 467.7317 (2010), pp. 820–823.
- [78] M. T. Hummon et al. ‘2D magneto-optical trapping of diatomic molecules’. In: *Physical review letters* 110.14 (2013), p. 143001.
- [79] S. Ding et al. ‘Sub-Doppler cooling and compressed trapping of YO molecules at  $\mu$  K temperatures’. In: *Physical Review X* 10.2 (2020), p. 021049.
- [80] V. Zhelyazkova et al. ‘Laser cooling and slowing of CaF molecules’. In: *Phys. Rev. A* 89 (5 2014), p. 053416. DOI: 10.1103/PhysRevA.89.053416.
- [81] B. Hemmerling et al. ‘Laser slowing of CaF molecules to near the capture velocity of a molecular MOT’. In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 49.17 (2016), p. 174001.
- [82] S. Truppe et al. ‘Molecules cooled below the Doppler limit’. In: *Nature Physics* 13.12 (2017), pp. 1173–1176.
- [83] J. Lim et al. ‘Laser cooled YbF molecules for measuring the electron’s electric dipole moment’. In: *Physical review letters* 120.12 (2018), p. 123201.
- [84] F. Schreck and K. v. Druten. ‘Laser cooling for quantum gases’. In: *Nature Physics* 17.12 (2021), pp. 1296–1304.
- [85] C. J. Foot. *Atomic physics*. Vol. 7. OUP Oxford, 2004.
- [86] D. J. McCarron et al. ‘Magnetic Trapping of an Ultracold Gas of Polar Molecules’. In: *Phys. Rev. Lett.* 121 (1 2018), p. 013202. DOI: 10.1103/PhysRevLett.121.013202.
- [87] N. B. Vilas et al. ‘Magneto-optical trapping and sub-Doppler cooling of a polyatomic molecule’. In: *Nature* 606.7912 (2022), pp. 70–74.
- [88] C. Hallas et al. ‘Optical Trapping of a Polyatomic Molecule in an  $\ell$ -Type Parity Doublet State’. In: *Phys. Rev. Lett.* 130 (15 2023), p. 153202. DOI: 10.1103/PhysRevLett.130.153202.

- [89] H. J. Williams et al. 'Characteristics of a magneto-optical trap of molecules'. In: *New Journal of Physics* 19.11 (2017), p. 113035. DOI: 10.1088/1367-2630/aa8e52.
- [90] L. Anderegg et al. 'Laser cooling of optically trapped molecules'. In: *Nature Physics* 14.9 (2018), pp. 890–893.
- [91] J. J. Bureau et al. 'Blue-Detuned Magneto-optical Trap of Molecules'. In: *Phys. Rev. Lett.* 130 (19 2023), p. 193401. DOI: 10.1103/PhysRevLett.130.193401.
- [92] H. Busch et al. 'Trap loss in a dual-species Rb - Ar\* magneto-optical trap'. In: *Physical Review A* 73.2 (2006), p. 023406.
- [93] K. BUTLER et al. 'A dual species MOT of Yb and Cs'. PhD thesis. University of Durham, 2014.
- [94] S. Sutradhar et al. 'Fast loaded dual species magneto optical trap of cold sodium and potassium atoms with light-assisted inter-species interaction'. In: *AIP Advances* 13.6 (June 2023), p. 065317. ISSN: 2158-3226. DOI: 10.1063/5.0154985.
- [95] S. Jurgilas et al. 'Collisions in a dual-species magneto-optical trap of molecules and atoms'. In: *New Journal of Physics* 23.7 (2021), p. 075004.
- [96] D. McCarron. 'Laser cooling and trapping molecules'. In: 51.21 (2018), p. 212001. DOI: 10.1088/1361-6455/aadfb.
- [97] M. Tarbutt. 'Laser cooling of molecules'. In: *Contemporary Physics* (2019).
- [98] M. Zeppenfeld et al. 'Optoelectrical cooling of polar molecules'. In: *Physical Review A* 80.4 (2009), p. 041401.
- [99] M. Zeppenfeld et al. 'Sisyphus cooling of electrically trapped polyatomic molecules'. In: *Nature* 491.7425 (2012), pp. 570–573.
- [100] N. J. Fitch and M. R. Tarbutt. 'Principles and Design of a Zeeman-Sisyphus Decelerator for Molecular Beams'. In: *ChemPhysChem* 17.22 (2016), pp. 3609–3623. DOI: 10.1002/cphc.201600656.
- [101] B. L. Augenbraun et al. 'Zeeman-Sisyphus Deceleration of Molecular Beams'. In: *Phys. Rev. Lett.* 127 (26 2021), p. 263002. DOI: 10.1103/PhysRevLett.127.263002.
- [102] H. Sawaoka et al. 'Zeeman-Sisyphus deceleration for heavy molecules with perturbed excited-state structure'. In: *Phys. Rev. A* 107 (2 2023), p. 022810. DOI: 10.1103/PhysRevA.107.022810.
- [103] K. M. Jones et al. 'Ultracold photoassociation spectroscopy: Long-range molecules and atomic scattering'. In: *Rev. Mod. Phys.* 78 (2 2006), pp. 483–535. DOI: 10.1103/RevModPhys.78.483.
- [104] S. Azizi, M. Aymar and O. Dulieu. 'Prospects for the formation of ultracold ground state polar molecules from mixed alkali atom pairs'. In: *The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics* 31.2 (2004), pp. 195–203. DOI: 10.1140/epjd/e2004-00159-2.
- [105] A. Ciamei et al. 'Efficient production of long-lived ultracold Sr<sub>2</sub> molecules'. In: *Phys. Rev. A* 96 (1 2017), p. 013406. DOI: 10.1103/PhysRevA.96.013406.
- [106] K. Leung et al. 'Ultracold 88Sr<sub>2</sub> molecules in the absolute ground state'. In: *New Journal of Physics* 23.11 (2021), p. 115002.

- [107] Y. Yu et al. ‘Coherent Optical Creation of a Single Molecule’. In: *Phys. Rev. X* 11 (3 2021), p. 031061. DOI: 10.1103/PhysRevX.11.031061.
- [108] T. Franzen et al. ‘Intercombination-line photoassociation spectroscopy of  $^{87}\text{Rb}^{170}\text{Yb}$ ’. In: *Phys. Rev. A* 107 (2 2023), p. 023114. DOI: 10.1103/PhysRevA.107.023114.
- [109] J. Schnabel et al. ‘Towards photoassociation processes of ultracold rubidium trimers’. In: *Phys. Rev. A* 103 (2 2021), p. 022820. DOI: 10.1103/PhysRevA.103.022820.
- [110] A. A. Elkamshishy and C. H. Greene. ‘Triatomic Photoassociation in an Ultracold Atom–Molecule Collision’. In: *The Journal of Physical Chemistry A* 127.1 (2022), pp. 18–28.
- [111] M.-S. Heo et al. ‘Formation of ultracold fermionic NaLi Feshbach molecules’. In: *Physical Review A* 86.2 (2012), p. 021602.
- [112] F. Wang et al. ‘Formation of ultracold NaRb Feshbach molecules’. In: *New Journal of Physics* 17.3 (2015), p. 035003.
- [113] C. Chin et al. ‘Feshbach resonances in ultracold gases’. In: *Rev. Mod. Phys.* 82 (2 2010), pp. 1225–1286. DOI: 10.1103/RevModPhys.82.1225.
- [114] K. Bergmann. ‘STIRAP: a historical perspective and some news’. In: *Molecular Beams in Physics and Chemistry: From Otto Stern’s Pioneering Exploits to Present-Day Feats*. Springer International Publishing Cham, 2021, pp. 445–462.
- [115] C.-H. Wu et al. ‘Ultracold Fermionic Feshbach Molecules of  $^{23}\text{Na}^{40}\text{K}$ ’. In: *Phys. Rev. Lett.* 109 (8 2012), p. 085301. DOI: 10.1103/PhysRevLett.109.085301.
- [116] M.-S. Heo et al. ‘Formation of ultracold fermionic NaLi Feshbach molecules’. In: *Phys. Rev. A* 86 (2 2012), p. 021602. DOI: 10.1103/PhysRevA.86.021602.
- [117] T. M. Rvachov et al. ‘Long-Lived Ultracold Molecules with Electric and Magnetic Dipole Moments’. In: *Phys. Rev. Lett.* 119 (14 2017), p. 143001. DOI: 10.1103/PhysRevLett.119.143001.
- [118] J. T. Zhang et al. ‘Forming a Single Molecule by Magnetoassociation in an Optical Tweezer’. In: *Phys. Rev. Lett.* 124 (25 2020), p. 253401. DOI: 10.1103/PhysRevLett.124.253401.
- [119] T. Takekoshi et al. ‘Ultracold dense samples of dipolar RbCs molecules in the rovibrational and hyperfine ground state’. In: *Physical review letters* 113.20 (2014), p. 205301.
- [120] A. Green et al. ‘Feshbach Resonances in  $p$ -Wave Three-Body Recombination within Fermi-Fermi Mixtures of Open-Shell  $^6\text{Li}$  and Closed-Shell  $^{173}\text{Yb}$  Atoms’. In: *Phys. Rev. X* 10 (3 2020), p. 031037. DOI: 10.1103/PhysRevX.10.031037.
- [121] T. Franzen et al. ‘Observation of magnetic Feshbach resonances between Cs and  $^{173}\text{Yb}$ ’. In: *Phys. Rev. Res.* 4 (4 2022), p. 043072. DOI: 10.1103/PhysRevResearch.4.043072.
- [122] B. Mukherjee, M. D. Frye and J. M. Hutson. ‘Magnetic Feshbach resonances between atoms in  $^2\text{S}$  and  $^3\text{P}_0$  states: Mechanisms and dependence on atomic properties’. In: *Phys. Rev. Res.* 5 (1 2023), p. 013102. DOI: 10.1103/PhysRevResearch.5.013102.
- [123] R. C. Bird, M. R. Tarbutt and J. M. Hutson. ‘Tunable Feshbach resonances in collisions of ultracold molecules in  $^2\Sigma$  states with alkali-metal atoms’. In: *Phys. Rev. Res.* 5 (2 2023), p. 023184. DOI: 10.1103/PhysRevResearch.5.023184.

- [124] D. K. Ruttley et al. 'Formation of Ultracold Molecules by Merging Optical Tweezers'. In: *Phys. Rev. Lett.* 130 (22 2023), p. 223401. DOI: 10.1103/PhysRevLett.130.223401.
- [125] D. C. Alexandre Trottier and E. Wrede. 'Photostop: production of zero-velocity molecules by photodissociation in a molecular beam'. In: *Molecular Physics* 109.5 (2011), pp. 725–733. DOI: 10.1080/00268976.2010.550142.
- [126] S. Matthews, S. Willitsch and T. Softley. 'Fully state-selected VMI study of the near-threshold photodissociation of NO 2: variation of the angular anisotropy parameter'. In: *Physical Chemistry Chemical Physics* 9.42 (2007), pp. 5656–5663.
- [127] B. Suk Zhao et al. 'Slow molecules produced by photodissociation'. In: *Journal of the Physical Society of Japan* 78.9 (2009), p. 094302.
- [128] A. Trottier, D. Carty and E. Wrede. 'Photostop: production of zero-velocity molecules by photodissociation in a molecular beam'. In: *Molecular Physics* 109.5 (2011), pp. 725–733.
- [129] W. Doherty et al. 'Production of cold bromine atoms at zero mean velocity by photodissociation'. In: *Physical Chemistry Chemical Physics* 13.18 (2011), pp. 8441–8447.
- [130] A. Rowland. 'Photostop'. PhD thesis. University of Durham, 2013.
- [131] J. S. Eardley et al. 'Magnetic trapping of SH radicals'. In: *Physical Chemistry Chemical Physics* 19.12 (2017), pp. 8423–8427.
- [132] J. EARDLEY. 'Trapping molecules using photostop'. PhD thesis. University of Durham, 2019.
- [133] T. Steimle and W. Ubachs. 'Introduction to the Special issue Spectroscopic Tests of Fundamental Physics'. In: *Journal of Molecular Spectroscopy* 300 (2014). Spectroscopic Tests of Fundamental Physics, pp. 1–2. ISSN: 0022-2852. DOI: 10.1016/j.jms.2014.04.004.
- [134] H. Pauly. *Atom, molecule, and cluster beams I: Basic theory, production and detection of thermal energy beams*. Vol. 28. Springer Science & Business Media, 2012.
- [135] J. van Veldhoven, H. L. Bethlem and G. Meijer. 'ac Electric Trap for Ground-State Molecules'. In: *Phys. Rev. Lett.* 94 (8 2005), p. 083001. DOI: 10.1103/PhysRevLett.94.083001.
- [136] E. R. Hudson et al. 'Cold Molecule Spectroscopy for Constraining the Evolution of the Fine Structure Constant'. In: *Phys. Rev. Lett.* 96 (14 2006), p. 143004. DOI: 10.1103/PhysRevLett.96.143004.
- [137] L. Semeria et al. 'Molecular-beam resonance method with Zeeman-decelerated samples: Application to metastable helium molecules'. In: *Phys. Rev. A* 98 (6 2018), p. 062518. DOI: 10.1103/PhysRevA.98.062518.
- [138] P. Sandars. 'Enhancement factor for the electric dipole moment of the valence electron in an alkali atom'. In: *Physics Letters* 22.3 (1966), pp. 290–291.
- [139] D. M. Kara et al. 'Measurement of the electron's electric dipole moment using YbF molecules: methods and data analysis'. In: *New Journal of Physics* 14.10 (2012), p. 103051.
- [140] C. J. Ho et al. 'New techniques for a measurement of the electron's electric dipole moment'. In: *New Journal of Physics* 22.5 (2020), p. 053031. DOI: 10.1088/1367-2630/ab83d2.

- [141] W. B. Cairncross et al. 'Precision Measurement of the Electron's Electric Dipole Moment Using Trapped Molecular Ions'. In: *Phys. Rev. Lett.* 119 (15 2017), p. 153001. DOI: 10.1103/PhysRevLett.119.153001.
- [142] T. S. Roussy et al. 'An improved bound on the electron's electric dipole moment'. In: *Science* 381.6653 (2023), pp. 46–50. DOI: 10.1126/science.adg4084.
- [143] S. Chefdeville et al. 'Appearance of low energy resonances in CO–para-H<sub>2</sub> inelastic collisions'. In: *Physical Review Letters* 109.2 (2012), p. 023201.
- [144] I. R. Sims and I. W. Smith. 'Gas-phase reactions and energy transfer at very low temperatures'. In: *Annual review of physical chemistry* 46.1 (1995), pp. 109–138.
- [145] J. J. Gilijamse et al. 'Near-threshold inelastic collisions using molecular beams with a tunable velocity'. In: *Science* 313.5793 (2006), pp. 1617–1620.
- [146] R. P. Feynman. 'Simulating physics with computers'. In: *International Journal of Theoretical Physics* 21 (1981), pp. 6–7.
- [147] I. M. Georgescu, S. Ashhab and F. Nori. 'Quantum simulation'. In: *Reviews of Modern Physics* 86.1 (2014), p. 153.
- [148] G. Pupillo et al. 'Condensed matter physics with cold polar molecules'. In: *Cold Molecules*. CRC Press, 2009, pp. 453–502.
- [149] I. Buluta and F. Nori. 'Quantum simulators'. In: *Science* 326.5949 (2009), pp. 108–111.
- [150] S. Chicco et al. 'Proof-of-Concept Quantum Simulator Based on Molecular Spin Qudits'. In: *Journal of the American Chemical Society* (2023).
- [151] P. A. Walker. 'MT-MOT: a Hybrid Magnetic Trap / Magneto-Optical Trap'. MRes thesis. University of Durham, 2019.
- [152] R. Campargue. 'Progress in overexpanded supersonic jets and skimmed molecular beams in free-jet zones of silence'. In: *The Journal of Physical Chemistry* 88.20 (1984), pp. 4466–4474. ISSN: 0022-3654. DOI: 10.1021/j150664a004.
- [153] K. Luria, W. Christen and U. Even. 'Generation and Propagation of Intense Supersonic Beams'. In: *The Journal of Physical Chemistry A* 115.25 (2011), pp. 7362–7367. ISSN: 1089-5639. DOI: 10.1021/jp201342u.
- [154] M. R. Tarbutt et al. 'Preparation and manipulation of molecules for fundamental physics tests'. In: *Cold Molecules*. CRC Press, 2009, pp. 587–628.
- [155] J. P. Toennies and K. Winkelmann. 'Theoretical studies of highly expanded free jets: Influence of quantum effects and a realistic intermolecular potential'. In: *The Journal of Chemical Physics* 66.9 (1977), pp. 3965–3979. DOI: 10.1063/1.434448.
- [156] U. Even. "'The Even-Lavie valve as a source for high intensity supersonic beam'". In: *EPJ Techniques and Instrumentation* 2.1 (2015), p. 17. ISSN: 2195-7045. DOI: 10.1140/epjti/s40485-015-0027-5.
- [157] U. Even and Z. Zhang. 'Pulsed Supersonic Beams from High Pressure Source: Simulation Results and Experimental Measurements'. In: *Advances in Chemistry* 636042 (2014).
- [158] K. Luria, N. Lavie and U. Even. 'Dielectric barrier discharge source for supersonic beams'. In: *Review of Scientific Instruments* 80.104102 (2009). DOI: 10.1063/1.3244085. eprint: 10.1063/1.3244085.

- [159] T. Halfmann, J. Koensgen and K. Bergmann. ‘A source for a high-intensity pulsed beam of metastable helium atoms’. In: *Measurement Science and Technology* 11.10 (2000), p. 1510.
- [160] J. C. Maxwell. *The Scientific Papers of James Clerk Maxwell*. Ed. by W. D. Niven. Vol. 2. Cambridge Library Collection - Physical Sciences. Cambridge University Press, 2011. DOI: 10.1017/CBO9780511710377.
- [161] A. J. Kox. ‘The discovery of the electron: II. The Zeeman effect’. In: *European Journal of Physics* 18.3 (1997), p. 139. DOI: 10.1088/0143-0807/18/3/003.
- [162] P. Zeeman. ‘Over de invloed eener magnetisatie op den aard van het door een stof uitgezonden licht.’ In: *Verslagen en Mededeelingen der Kon. Academie van Wetenschappen, Afd. Natuurkunde* 5 (Jan. 1896), pp. 181–184.
- [163] T. Preston. ‘The Zeeman Effect Photographed’. In: *Nature* 57.1469 (1897), pp. 173–173.
- [164] T. Preston. ‘Radiation phenomena in a strong magnetic field:...’ In: R. Dublin Soc. 1898.
- [165] A. H. Compton. ‘The magnetic electron’. In: *Journal of the Franklin Institute* 192.2 (1921), pp. 145–155. ISSN: 0016-0032. DOI: 10.1016/S0016-0032(21)90917-7.
- [166] W. Pauli. ‘Über den Zusammenhang des Abschlusses der Elektronengruppen im Atom mit der Komplexstruktur der Spektren’. In: *Einführung und Originaltexte* (1925), p. 229.
- [167] J. Toscano et al. ‘Evolutionary algorithm optimization of Zeeman deceleration: Is it worthwhile for longer decelerators?’ In: *The Journal of Physical Chemistry A* 123.25 (2019), pp. 5388–5394.
- [168] T. Damjanović et al. ‘A new design for a traveling-wave Zeeman decelerator: II. Experiment’. In: *New Journal of Physics* 23.10 (2021), p. 105007. DOI: 10.1088/1367-2630/ac2c2b.
- [169] E. Narevicius et al. ‘Towards magnetic slowing of atoms and molecules’. In: *New Journal of Physics* 9.4 (2007), p. 96. DOI: 10.1088/1367-2630/9/4/096.
- [170] J. Narevicius and E. Narevicius. ‘Multichannel high peak power tunable duration pulse generation for the moving magnetic trap decelerator’. In: *Review of scientific instruments* 93.1 (2022).
- [171] P. Elleaume, O. Chubar and J. Chavanne. ‘Computing 3D magnetic fields from insertion devices’. In: *Proceedings of the 1997 Particle Accelerator Conference (Cat. No.97CH36167)* 3.4 (1998), pp. 3509–3511. DOI: 10.1109/PAC.1997.753258.
- [172] D. C. Meeker. *Finite Element Method Magnetics, Version 4.2 (28Feb2018 Build)*. URL: <http://www.femm.info>.
- [173] J. B. Anderson and J. B. Fenn. ‘Velocity Distributions in Molecular Beams from Nozzle Sources’. In: *The Physics of Fluids* 8.5 (1965), pp. 780–787. DOI: 10.1063/1.1761320.
- [174] D. R. Miller and R. P. Andres. ‘Rotational Relaxation of Molecular Nitrogen’. In: *The Journal of Chemical Physics* 46.9 (1967), pp. 3418–3423. DOI: 10.1063/1.1841233.
- [175] P. N. Bajaj and P. K. Chakraborti. ‘Characterization of supersonic beams by time-of-flight technique’. In: *Pramana* 38.4 (Apr. 1992), pp. 397–409. DOI: 10.1007/BF02875387.
- [176] A. H. Futch and F. A. Grant. ‘Mean Life of the  $^3\text{P}_2$  Metastable Argon Level’. In: *Physical Review* 104.2 (1956), pp. 356–361. DOI: 10.1103/PhysRev.104.356.

- [177] G. Lochead. ‘Excited state spatial distributions in a cold strontium gas’. PhD thesis. University of Durham, 2012.
- [178] I. Hughes and T. Hase. *Measurements and their Uncertainties: A practical guide to modern error analysis*. OUP Oxford, 2010. ISBN: 9780199566334.
- [179] P. Virtanen et al. ‘SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python’. In: *Nature Methods* 17 (2020), pp. 261–272. DOI: 10.1038/s41592-019-0686-2.
- [180] G. E. P. Box. ‘Science and Statistics’. In: *Journal of the American Statistical Association* 71.356 (1976), pp. 791–799. DOI: 10.1080/01621459.1976.10480949.
- [181] C. R. Harris et al. ‘Array programming with NumPy’. In: *Nature* 585.7825 (Sept. 2020), pp. 357–362. DOI: 10.1038/s41586-020-2649-2.
- [182] M. Matsumoto and T. Nishimura. ‘Mersenne Twister: A 623-Dimensionally Equidistributed Uniform Pseudo-Random Number Generator’. In: 8.1 (1998), pp. 3–30. ISSN: 1049-3301. DOI: 10.1145/272991.272995.
- [183] M. E. O’Neill. *PCG: A Family of Simple Fast Space-Efficient Statistically Good Algorithms for Random Number Generation*. Tech. rep. HMC-CS-2014-0905. Claremont, CA: Harvey Mudd College, 2014.
- [184] L. Devroye. ‘Non-Uniform Random Variate Generation’. In: 1986.
- [185] G. Casella, C. P. Robert and M. T. Wells. ‘Generalized Accept-Reject Sampling Schemes’. In: *Lecture Notes-Monograph Series* 45 (2004), pp. 342–347. ISSN: 07492170.
- [186] L. Verlet. ‘Computer ‘Experiments’ on Classical Fluids. I. Thermodynamical Properties of Lennard-Jones Molecules’. In: *Phys. Rev.* 159 (1 July 1967), pp. 98–103. DOI: 10.1103/PhysRev.159.98.
- [187] G. Woodgate. *Elementary atomic structure, clarendon*. 1980.
- [188] G. Malli and C. Froese. ‘Fermi contact interaction constants calculated from numerical hartree-fock wave functions’. In: *International Journal of Quantum Chemistry* 1.S1 (1967), pp. 111–114.
- [189] S. Booshanam. ‘Simulation of sympathetic cooling of hydrogen with laser cooled lithium in an MT-MOT’. MPhys thesis. University of Durham, 2021.
- [190] S. Freake. *Electromagnetic Fields*. Open University. Course SMT359. Open University Worldwide, 2006. ISBN: 9780749269869.
- [191] C. Shannon. ‘Communication in the Presence of Noise’. In: *Proceedings of the IRE* 37.1 (1949), pp. 10–21. DOI: 10.1109/JRPROC.1949.232969.
- [192] C. Runge. ‘Über empirische Funktionen und die Interpolation zwischen äquidistanten Ordinaten’. In: *Zeitschrift für Mathematik und Physik* 46 (1901), pp. 224–243.
- [193] J. H. Ahlberg, E. N. Nilson and J. L. Walsh. *The Theory of Splines and Their Applications: Mathematics in Science and Engineering: A Series of Monographs and Textbooks, Vol. 38*. Vol. 38. Elsevier, 1967, pp. 1–8.
- [194] I. J. Schoenberg. ‘Contributions to the Problem of Approximation of Equidistant Data by Analytic Functions’. In: *I. J. Schoenberg Selected Papers*. Ed. by C. de Boor. Boston, MA: Birkhäuser Boston, 1988, pp. 58–87. ISBN: 978-1-4899-0433-1. DOI: 10.1007/978-1-4899-0433-1\_2.

- [195] G. Birkhoff and C. de Boor. ‘Piecewise polynomial interpolation and approximation’. In: *Proc. General Motors Symposium of 1964*. Ed. by H. L. Garabedian. Amsterdam: Elsevier, 1965, pp. 164–190.
- [196] C. A. Hall and W. Meyer. ‘Optimal error bounds for cubic spline interpolation’. In: *Journal of Approximation Theory* 16.2 (1976), pp. 105–122. ISSN: 0021-9045.
- [197] F. Lekien, C. Coulliette and J. Marsden. ‘Lagrangian Structures in Very High-Frequency Radar Data and Optimal Pollution Timing’. In: *American Institute of Physics Conference Proceedings* 676 (Aug. 2003). DOI: 10.1063/1.1612209.
- [198] F. Lekien and J. Marsden. ‘Tricubic interpolation in three dimensions’. In: *International Journal for Numerical Methods in Engineering* 63.3 (2005), pp. 455–471. DOI: 10.1002/nme.1296.
- [199] E. Catmull and R. Rom. ‘A class of local interpolating splines’. In: *Computer Aided Geometric Design*. Ed. by R. Barnhill and R. Riesenfeld. Academic Press, 1974, pp. 317–326. ISBN: 978-0-12-079050-0. DOI: 10.1016/B978-0-12-079050-0.50020-5.
- [200] W. H. Press et al. *Numerical recipes in C++ : the art of scientific computing*. 2002.
- [201] K. Bradach. ‘A Triquintic Interpolation Method with Global C2 Continuity’. MA thesis. San Francisco State University, 2022.
- [202] P.-b. Zhou. *Numerical Analysis of Electromagnetic Fields*. Springer, 1993.
- [203] M. Chilenski, I. Faust and J. Walk. ‘eqtools: Modular, extensible, open-source, cross-machine Python tools for working with magnetic equilibria’. In: *Computer Physics Communications* 210 (Sept. 2016), pp. 155–162. DOI: 10.1016/j.cpc.2016.09.011.
- [204] G. Van Rossum and F. L. Drake. *Python 3 Reference Manual*. Scotts Valley, CA: CreateSpace, 2009. ISBN: 1441412697.
- [205] P. A. Walker, U. Krohn and D. Carty. ‘ARBTools: A Tricubic Spline Interpolator for Three-Dimensional Scalar or Vector Fields’. In: *Journal of Open Research Software* (2019). DOI: 10.5334/jors.258.
- [206] P. A. Walker. *Quadcubic interpolation: a four-dimensional spline method*. 2019. arXiv: 1904.09869 [math.NA].
- [207] K. Franke et al. ‘Behavioral state tunes mouse vision to ethological features through pupil dilation’. In: *bioRxiv* (2021). DOI: 10.1101/2021.09.03.458870.
- [208] L. Gerster et al. ‘Experimental Bayesian Calibration of Trapped-Ion Entangling Operations’. In: *PRX Quantum* 3 (2 2022), p. 020350. DOI: 10.1103/PRXQuantum.3.020350.
- [209] S. S. Tie et al. ‘Constraining IGM enrichment and metallicity with the C iv forest correlation function’. In: *Monthly Notices of the Royal Astronomical Society* 515.3 (Aug. 2022), pp. 3656–3673. ISSN: 0035-8711. DOI: 10.1093/mnras/stac2021.
- [210] M. D. Louwerse and D. A. Sivak. ‘Multidimensional minimum-work control of a 2D Ising model’. In: *The Journal of Chemical Physics* 156.19 (2022), p. 194108. DOI: 10.1063/5.0086079.
- [211] J. D. Jackson and L. C. Levitt. ‘Classical Electrodynamics’. In: *Physics Today* 15.11 (1962), p. 62. DOI: 10.1063/1.3057859.

- [212] W. Press et al. *Numerical Recipes: The Art of Scientific Computing (3rd ed.)* Cambridge University Press, 2007.
- [213] M. P. Allen and D. J. Tildesley. ‘Point Location’. In: *Computer Simulation of Liquids, Second Edition*. Oxford University Press, 2017. ISBN: 9780198803195. DOI: 10.1093/oso/9780198803195.001.0001.
- [214] D. Frenkel and B. Smit. In: *Understanding Molecular Simulation: From Algorithms to Applications*. Orlando: Academic Press, Inc., 1996.
- [215] E. De Lara, A. LaMarca and M. Satyanarayanan. In: *Location Systems: An Introduction to the Technology Behind Location Awareness*. Morgan & Claypool Publishers, 2008.
- [216] M. de Berg et al. ‘Point Location’. In: *Computational Geometry: Algorithms and Applications*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1997, pp. 119–144. ISBN: 978-3-662-03427-9. DOI: 10.1007/978-3-662-03427-9\_6.
- [217] B. Bertsche, J. Jankunas and A. Osterwalder. ‘Low-temperature Collisions between Neutral Molecules in Merged Molecular Beams’. In: *CHIMIA* 68.4 (2014), p. 256. DOI: 10.2533/chimia.2014.256.
- [218] E. M. McMillan. ‘The Synchrotron—A Proposed High Energy Particle Accelerator’. In: *Phys. Rev.* 68 (5-6 1945), pp. 143–144. DOI: 10.1103/PhysRev.68.143.
- [219] J. J. Gilijamse et al. ‘Optimizing the Stark-decelerator beamline for the trapping of cold molecules using evolutionary strategies’. In: *Phys. Rev. A* 73 (6 2006), p. 063410. DOI: 10.1103/PhysRevA.73.063410.
- [220] A. W. Wiederkehr, S. D. Hogan and F. Merkt. ‘Phase stability in a multistage Zeeman decelerator’. In: *Phys. Rev. A* 82 (4 2010), p. 043428. DOI: 10.1103/PhysRevA.82.043428.
- [221] M. N. Bera. ‘Zeeman Deceleration of Supersonic Beam trapping of Paramagnetic Atoms in a Traveling Magnetic Wave’. PhD thesis. Université Paris Sud, 2011.
- [222] A. Trimeche. ‘Décélération Zeeman-Stern Gerlach d’un jet supersonique de particules paramagnétiques par une onde de champ magnétique progressive’. PhD thesis. Université Paris Sud, 2014.
- [223] K. Maeda, M. L. Wall and L. D. Carr. ‘Hyperfine structure of the hydroxyl free radical (OH) in electric and magnetic fields’. In: *New Journal of Physics* 17.4 (2015), p. 045014. DOI: 10.1088/1367-2630/17/4/045014.
- [224] B. Lynch. ‘Making slow beams from photostop and Zeeman deceleration’. MPhys thesis. University of Durham, 2019.
- [225] D. Nohlman. ‘A Permanent Magnet Trap For Buffer Gas Cooled Atoms’. PhD thesis. Imperial College London, 2015.
- [226] D. DeMille, D. R. Glenn and J. Petricka. ‘Microwave traps for cold polar molecules’. In: *The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics* 31 (2 2004), pp. 375–384. DOI: 10.1140/epjd/e2004-00163-6.
- [227] A. O. G. Wallis and J. M. Hutson. ‘Production of Ultracold NH Molecules by Sympathetic Cooling with Mg’. In: *Phys. Rev. Lett.* 103 (18 2009), p. 183201. DOI: 10.1103/PhysRevLett.103.183201.

- [228] M. Lara et al. ‘Ultracold Rb-OH Collisions and Prospects for Sympathetic Cooling’. In: *Phys. Rev. Lett.* 97 (18 2006), p. 183201. DOI: 10.1103/PhysRevLett.97.183201.
- [229] M. H. Alexander. ‘Rotationally inelastic collisions between a diatomic molecule in a  $^2\Sigma^+$  electronic state and a structureless target’. In: *The Journal of Chemical Physics* 76.7 (Apr. 1982), pp. 3637–3645. ISSN: 0021-9606. DOI: 10.1063/1.443401.
- [230] R. V. Krems et al. ‘Spin-flipping transitions in  $^2\Sigma$  molecules induced by collisions with structureless atoms’. In: *Phys. Rev. A* 67 (6 2003), p. 060703. DOI: 10.1103/PhysRevA.67.060703.
- [231] M. Lara et al. ‘Cold collisions between OH and Rb: The field-free case’. In: *Physical Review A* 75.1 (2007), p. 012704.
- [232] M. L. González-Martínez and J. M. Hutson. ‘Ultracold hydrogen atoms: a versatile coolant to produce ultracold molecules’. In: *Physical review letters* 111.20 (2013), p. 203004.
- [233] A. O. Wallis et al. ‘The prospects of sympathetic cooling of NH molecules with Li atoms’. In: *The European Physical Journal D* 65 (2011), pp. 151–160.
- [234] P. Soldán and J. M. Hutson. ‘Interaction of  $\text{NH}(X^3\Sigma^-)$  Molecules with Rubidium Atoms: Implications for Sympathetic Cooling and the Formation of Extremely Polar Molecules’. In: *Phys. Rev. Lett.* 92 (16 2004), p. 163202. DOI: 10.1103/PhysRevLett.92.163202.
- [235] T. V. Tscherbul and J. Kłos. ‘Magnetic tuning of ultracold barrierless chemical reactions’. In: *Phys. Rev. Res.* 2 (1 2020), p. 013117. DOI: 10.1103/PhysRevResearch.2.013117.
- [236] M. Thompson. ‘Cooling Molecules to Ultracold Temperatures’. MPhys thesis. University of Durham, 2018.
- [237] B. Boland et al. ‘Microwave spectroscopy of nonlinear free radicals III. High field Zeeman effect in HCO and DCO’. In: *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 360.1703 (1978), pp. 507–528.
- [238] R. Côté et al. ‘Enhanced cooling of hydrogen atoms by lithium atoms’. In: *Physical Review Letters* 84.13 (2000), p. 2806.
- [239] I. Setija et al. ‘Optical cooling of atomic hydrogen in a magnetic trap’. In: *Physical review letters* 70.15 (1993), p. 2257.
- [240] F. Gadéa, T. Leininger and A. Dickinson. ‘Accurate calculation of the scattering length for the cooling of hydrogen atoms by lithium atoms’. In: *The European Physical Journal D-Atomic, Molecular, Optical and Plasma Physics* 15 (2001), pp. 251–255.
- [241] O. Turnbull. ‘Simulating the Laser Cooling of Lithium Atoms’. MPhys thesis. University of Durham, 2020.
- [242] R. K. Hanley et al. ‘Quantitative simulation of a magneto-optical trap operating near the photon recoil limit’. In: *Journal of Modern Optics* 65.5-6 (2018), pp. 667–676.
- [243] J. Gilmour et al. ‘A resonance ionization mass spectrometer for xenon’. In: *Measurement Science and Technology* 2.7 (1991), p. 589.
- [244] S. D. Hogan et al. ‘Slow beams of atomic hydrogen by multistage Zeeman deceleration’. In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 41.8 (2008), p. 081005.
- [245] N. Bulleid et al. ‘Characterization of a cryogenic beam source for atoms and molecules’. In: *Physical Chemistry Chemical Physics* 15.29 (2013), pp. 12299–12307.

- [246] N. E. Bulleid. ‘Slow, cold beams of polar molecules for precision measurements’. PhD thesis. Imperial College London, 2013.
- [247] A. Messiah. *Quantum Mechanics*. (vol. 1. Translated by G.M. Temmer. vol. 2. Translated by J. Potter.) North Holland Publishing Co., 1961.
- [248] M. Lara, B. L. Lev and J. L. Bohn. ‘Loss of molecules in magneto-electrostatic traps due to nonadiabatic transitions’. In: *Phys. Rev. A* 78 (3 2008), p. 033433. DOI: 10.1103/PhysRevA.78.033433.