Cold state-selected radicals for the study of low temperature chemistry

Jutta Toscano

Merton College, University of Oxford

A thesis submitted for the degree of Doctor of Philosophy

Michaelmas Term, 2018

Abstract

Cold state-selected radicals for the study of low temperature chemistry

Jutta Toscano, Merton College

A thesis submitted in partial fulfilment of the requirements
for the degree of Doctor of Philosophy of the University of Oxford

Michaelmas Term, 2018

Generating a controllable and pure source of molecular free-radicals or open-shell atoms has been one of the primary barriers hindering the detailed study of radical processes in the laboratory. In this thesis, a novel source of state-selected radicals – composed of a Zeeman decelerator interfaced with a newly-designed magnetic guide – is introduced. This tuneable source generates a pure beam of velocity-selected hydrogen atoms that will enable the study of radical interactions with exceptional control over the properties of the radical species.

In Zeeman deceleration, time-varying spatially inhomogeneous magnetic fields are used to create packets of translationally cold, quantum-state-selected paramagnetic particles with a tuneable forward velocity, which are ideal for cold reaction dynamics studies. In a proof-of-principle experiment, the Zeeman deceleration of nitrogen atoms in the metastable $^2D_{5/2}$ state from 460 to $410\,\mathrm{m\,s^{-1}}$ is demonstrated for the first time. The covariance matrix adaptation evolutionary strategy (CMA-ES) is adopted in order to optimise deceleration switching sequences for the operation of a 12-stage Zeeman decelerator. Using the optimised sequences, a 40% increase in the number of decelerated $H(^2S_{1/2})$ atoms is observed compared to standard sequences for the same final velocity, imposing the same experimental boundary conditions. Furthermore, up to 98% of the initial kinetic energy of particles in the incoming beam is removed by the optimised sequences, compared to the removal of a maximum of 83% of kinetic energy with standard sequences. Three-dimensional particle-trajectory simulations show that the optimal solution uncovered by the evolutionary algorithm is not merely a local optimisation of the experimental parameters – it is a new mode of operation that goes beyond the standard periodic phase stability approach typically adopted.

A novel magnetic guide is designed and constructed to purify the post-deceleration beam. Only radicals with a selected velocity are transmitted through the guide; all other components of the incoming beam (radical species travelling at other velocities, precursor molecules and seed gases) are removed. The guide is composed of four Halbach arrays – hexapolar focusing elements – and two skimming blades. The relative positions of these components can be adjusted to tune the properties of the resulting beam and to optimise transmission for a given velocity. Experimental measurements of Zeeman-decelerated H atoms transmitted through the guide, combined with extensive simulations, show that the magnetic guide successfully removes 99% of H atoms travelling outside the narrow target velocity range.

Acknowledgements

I wish to express my gratitude towards the people who have contributed to the work presented in this thesis, have supported me throughout my DPhil and have made these four years enriching and fulfilling.

Tim Softley, for being an inspiring and kind supervisor who encouraged me with patience and wit.
Brianna Heazlewood, for being a supportive and caring supervisor with unwavering passion and enthusiasm.
Chris Rennick, for teaching me the arts of programming, designing, troubleshooting, fixing and persevering.
Katrin Dulitz, for entrusting me with the decelerator, teaching me all the tricks and always being available to help.
Lorenzo Petralia, for being like a brother to me throughout these years.
Andriana Tsikritea for her optimism, Michal Hejduk and Mike Kohlhoff for their pessimism.
All the Softley and Heazlewood group members, for challenging me by joining – with varying degrees of enthusiasm – the lunchtime debate club.
The PTCL mechanical and electronic workshops, for their assiduous help and support, and for their friendly smiles towards anyone who brings them more trouble. Especially Andrew Green and Les Hill, for building the beam deflector assembly and the magnetic guide, respectively, Neville Baker for making new solenoid coils and John Adams for repeatedly fixing our old YAG laser.
The CRL and PTCL facilities teams, especially Paul Mitchell and Martin Lyne for their tireless assistance during the lab move, as well as afterwards; their cheerful and kind temperament made all the difference.
All my Oxford friends, for the countless evenings spent together.
Michael Fox, for brightening my days all along.

Ringrazio i miei amici di sempre per essere parte della mia vita e farmi essere parte della loro. Piera Porciatti con la sua intuizione ed empatia, Irene Scarpelli con la sua saggezza e pazienza, Deianira Iozzi con la sua tenacia e curiosità, Marta De Benedictis con la sua spontaneità e spensieratezza, Antonio Morbidelli con il suo idealismo e altruismo, Dario Luzzi con la sua eloquenza e determinazione, Samuel Okpokpo con la sua allegria e il suo affetto, Fulvia Zamboni con la sua originalità e il suo umorismo. Ringrazio anche le loro famiglie per essere anche un po' famiglia mia.

Infine ringrazio la mia mamma, Clara Toscano, che è la persona più forte, indipendente e coraggiosa che conosca. Grazie per ispirarmi a seguire i miei sogni e per sorreggermi mentre ci provo. Questa tesi è dedicata a te – il merito ce lo smezziamo, come sempre.

Contents

Page

1 | Introduction

1.1 Cold chemistry

At room temperature, chemical reactions typically occur between reagents that are present in a variety of different quantum states and that possess a distribution of velocities. The kinetics and dynamics measured in this temperature regime reflect the average behaviour of the reacting partners which, albeit valuable to our macroscopic understanding of gas-phase reactions, precludes our detailed microscopic interpretation of it. Disentangling the contribution of the individual quanta of internal energy of the reactants on the reaction outcome at the atomic level is not possible under such conditions.

At lower temperatures, on the other hand, fewer quantum states are populated and the kinetic energy of the reactants is lower. While reactions with an appreciable activation energy barrier become slower as the temperature is decreased, barrier-less reactions involving at least one transient, highly-reactive species can instead become faster. Indeed, low temperature reactivity can be dominated by long-range inter-molecular interactions and its experimental study permits the detailed investigation of computationally intractable potential energy surfaces. Additionally, the wave nature of reactant particles becomes significant at low temperatures, where their de Broglie wavelength becomes larger than the length scales of chemical bonding, and quantum-mechanical effects such as tunnelling and resonances can be observed [1–5].

Investigating chemical processes in the cold regime (below 1 K) opens the possibility of studying perfectly controlled reactions, where the velocity and quantum state of the reactants is accurately known and the velocity and quantum state of the products can be probed. Once this input–output, state-to-state relationship is established and extensively understood, one could envisage manipulating the outcome of a reaction by precisely applying electromagnetic fields to deterministically obtain one product over another with 100% yield, gaining ultimate control over chemical reactions.

The study of low temperature reactivity is also important to understand the chemistry that occurs in the cold parts of the Universe, such as the interstellar medium and planetary atmospheres, where temperatures can be as low as 10 K. Diffuse and dense interstellar clouds are rich in chemicals and their chemistry is dominated by processes involving highly-reactive species such as ions and radicals, since reactions between electronically-stable species mostly have prohibitive activation energies at the low ambient temperatures. These transient species are abundant in such environments due to the relative ease with which they are formed, through photolysis and photo-ionisation by stellar UV radiation, and the reduced probability of being destroyed, given the very low gas densities which lead to infrequent collisions. Whereas atomic and diatomic species are prevalent in diffuse clouds – where the destructive photochemistry caused by cosmic radiation breaks up larger species – numerous polyatomic species have been detected in dense clouds, where a larger number of dust grains is present to absorb and scatter radiation, thus shielding the chemicals present at the core of the cloud [6–9].

Investigating the rates and temperature dependencies of the reactions that occur in these cold environments is necessary in order to accurately model their chemistry. Reliable models of the interstellar medium, in which the predicted abundances of all species agree with the experimentally observed abundances, can be used to understand and predict the evolution of different clouds, for instance to simulate stellar lifecycles, or even to investigate the possible development of prebiotic molecules from simple organic species, which could explain the origin of life [10–12].

Whilst reactions between one electronically-stable species and one highly-reactive species have been extensively studied under various temperature regimes, reactions between two highly-reactive species, such as those between ions and radicals, have scarcely been explored in comparison, especially with respect to the temperature dependence of the reaction rates [13–15]. This is partly due to the challenges associated with carrying out experiments with sufficient densities of the two highly-reactive species and partly because these reactions do not occur frequently under atmospheric conditions. Indeed, in a relatively high pressure regime where highly-reactive species are present in small amounts, they are statistically bound to encounter, and react with, an electronically-stable partner instead of another highly-reactive species. However, in cold environments such as the interstellar medium, ion–radical reactions are an integral part of the chemical cycle and the lack of experimental measurements of their reaction rates and temperature dependencies has been identified as one of the weaknesses of current theoretical models of these complex environments [16, 17].

1.1.1 Ion–radical reactions

The majority of ion–radical reactions of interstellar relevance that have been experimentally studied so far have been investigated using two distinct techniques [14]. The first technique, selected ion flow tube (SIFT), relies on the injection of ions and radicals into a communal beam where they react; the second technique, atomic beam 22-pole ion trap (AB-22PT), consists of a beam of radicals directed into a static ionic target.

In the SIFT apparatus [14, 18], ions are generated (by electron impact ionisation, for instance) and mass selected with a quadrupole mass filter before being injected into a room-temperature flow of helium buffer gas, which collisionally thermalises the translational and internal degrees of freedom of the ions. Radicals are generated (by thermal dissociation or microwave discharge) within a flow of precursor molecules before being injected into the helium beam, downstream with respect to the ions. The ionic products of the reactions occurring within the flow are mass analysed and detected by an electron multiplier at the end of the flow tube. The decrease of ionic reactants and increase of ionic products are thus monitored as a function of reactant flow rate and reaction time, and they are compared to the response obtained when no radicals are present (i.e. when the radical-generating source is off). The experimental results are simulated using kinetic models to determine the contribution of the various reaction channels that are available for each system.

In the AB-22PT apparatus [19, 20], ions are generated by electron impact in a storage ion source, and they are injected into the 22-pole ion trap through an electrostatic quadrupole mass selector. In the trap, which can be cryogenically cooled to temperatures as low as 15 K, the ions are confined in space by radiofrequency fields and they are internally cooled by collisions with helium buffer gas. Radicals, specifically hydrogen atoms, are generated within a beam of precursor hydrogen molecules by a radio frequency discharge. The effusive H/H_2 beam flies through a copper tube where it thermalises to the temperature of the vessel (10 to 300 K), before being differentially pumped twice and focused by hexapolar magnets towards the ion trap entrance. Once the radicals enter the trap, reactions occur and the ionic content of the trap is probed (i.e. ions are ejected and analysed by mass spectrometry) after variable time intervals from the injection of the radicals. Very slow reactions can be measured as a result of the long trapping times achievable of milliseconds to minutes. The response obtained when the radicals are present (when the discharge is on) is compared to when the radicals are absent (discharge off). Kinetic models are then used to unravel the extent of the various possible reactions in order to rationalise the observed ionic abundances.

Both methods have proven to be extremely valuable tools, establishing reaction rate constants for a wealth of ion–radical reactions. However, they are both afflicted by relatively high uncertainties due to several factors, including the possible generation of metastable species, the recombination of radicals, the difficulty in determining the radical number density and the presence of precursor molecules, which sometimes react with the ions faster than the radicals do. For instance, the rate constant for the reaction between H atoms and CO_2^+ ions was determined to be $4.7 \times 10^{-10}\,\mathrm{cm^3 s^{-1}}$ (±30%) at room temperature by SIFT [21] and $4.5 \times 10^{-10}\,\mathrm{cm^3 s^{-1}}$ (±40%) by AB-22PT [19], where the ions were kept at room temperature whilst the H-atom beam was thermalised to 55 K. In both cases, the competing reaction of CO_2^+ with H_2 was also occurring, at a faster rate with a rate constant of $8.7 \times 10^{-10}\,\mathrm{cm^3 s^{-1}}$ (±15%) [21], resulting in the relatively high uncertainties.

Although measuring rate constants for thermally equilibrated systems, such as the ones described above, affords the advantage of precisely known populations of states and velocity distributions, measuring the reactivity of individual quantum states offers the possibility of observing quantum effects and shedding light on fundamental state-to-state dynamics over a range of translational energies and perhaps even as a function of relative orientation of the reactants [22]. Methods to reliably manipulate and control reactants are therefore necessary to carry out experiments with single-quantum-state selectivity. In the following section, a selection of techniques that enable the manipulation of ions and of radicals is introduced.

1.2 Controlling reactants

1.2.1 Ion manipulation

Several techniques have been developed to study the bimolecular reactions of ions, predominantly with electronically-stable species. Flow methods include the SIFT technique described in section 1.1.1, where the ions are injected into an inert gas flow, with which they thermalise before reacting. The walls of the flow tube can be cooled using a coolant or cryogen to decrease the temperature of the flow, but condensation of the reactants onto the tube surface becomes a problem at low temperatures where their vapour pressure becomes small. Alternatively, cooling of the flow can be attained by supersonically expanding gas through a Laval nozzle, which is the working principle of the CRESU technique (*cinétique de réaction en ecoulement supersonique uniforme* – reaction kinetics in uniform supersonic flow) [23]. The Laval nozzle is composed of a convergent section, where the gas is accelerated towards a lower pressure chamber without being compressed, and a divergent section, where it expands adiabatically generating a collimated dense flow with reduced temperature and density but increased velocity compared to the gas in the reservoir. The uniform, thermally equilibrated flow obtained is comparable to that of a flow tube without walls, and therefore without condensation of the reactant species present within it. Numerous reactions have been investigated by CRESU at temperatures as low as 5.8 K [24], including reactions between electronically-stable species and ions or radicals. Many of these reactions involving an unstable reactant – such as the reactions of N^+ with H_2O [25] or NH_3 [26] for example – have been observed to proceed faster as the temperature is lowered [16]. One limitation that is common to all flow methods is the intrinsic restriction to measuring relatively fast reactions, since the flow moves quickly and the flow tube is finite in length.

Trap methods, on the contrary, afford long interaction times by making use of the ease with which charged particles can be manipulated and confined in space using electric fields. In the AB-22PT apparatus described in section 1.1.1, the radiofrequency (RF) trap is composed of 22 cylindrical electrodes distributed around a circle. Oscillating voltages are applied to the electrodes such that adjacent rods have opposite polarities at all times and the resulting time-varying, inhomogeneous electric field forces the ions to oscillate around the centre of the trap [27]. The confined ions are then cooled by collision with a buffer gas such as helium, provided that the buffer gas does not react with the ions and the elastic collision cross section is large enough for thermalisation to be efficient. Another commonly used ion trap is the linear Paul (quadrupole) trap, which consists of four cylindrical electrodes (rather than 22) but works on a similar principle to the 22-pole one. Fewer electrodes make the effective trap volume of a linear Paul trap smaller but, at the same time, increase the optical access to the trap centre, which is advantageous for laser cooling the ions down to milliKelvin temperatures [3].

Certain atomic ions, such as the alkaline-earth ions Be^+, Mg^+, Ca^+ and Ba^+, can be laser cooled by repeated cycles of stimulated absorption and spontaneous emission of photons. In Doppler cooling, the ions travelling towards a laser beam that has a frequency slightly below that of an electronic transition may absorb a photon. When the excited ions fluoresce back to the ground state, they emit a photon isotropically such that, over numerous cycles, there is net momentum transfer to the ions in the opposite direction to their propagation, and they are slowed down. Below a critical temperature, the laser-cooled, trapped ions can form a Coulomb crystal, a structure where the gas-phase ions, although still mobile, tend to occupy specific positions in space (with ion–ion distance in the order of tens of µm) and resemble the structure of a solid crystal [28, 29]. The constant fluorescence of the ions induced by the continuous laser cooling allows the direct observation of these Coulomb crystal structures within the trap using a charge-coupled device (CCD) camera. Whilst molecular ions cannot usually be laser cooled due to their complex energy level structures, they can be loaded into a Coulomb crystal and sympathetically cooled by the laser-cooled ions that act as a heat sink. Sympathetic cooling only reduces the translational energy of the molecular ions; their internal energy is not affected unless a buffer gas is also present in the trap [30]. Depending on whether the molecular ions are heavier or lighter than the laser-cooled ions, they assemble on the outside or at the centre of the Coulomb crystal respectively and, although they do not fluoresce, their presence is inferred by the shape of the fluorescing laser-cooled ions in the crystal. Reactions with incoming particles can be studied by monitoring changes in the Coulomb crystal as a function of time: the identity of the species present in the crystal is established by mass spectrometry and an imaging algorithm is used to match the observed fluorescence patterns in order to extrapolate the number of ions of each species present. The high sensitivity of this technique makes Coulomb crystals an ideal target to study cold chemical reactions involving ions [28, 29, 31].

Coulomb crystals have been employed to investigate reactions of ions (either the laser-cooled ions or the sympathetically cooled ions) with several neutral (non-radical) species, including $H_2 + Mg^+$ [32], $HD + Be^+$ [33] and $ND_3 + Ca^+$ [34] among others [31]. For instance, Coulomb crystals of Ca^+ ions have been reacted with velocity-selected CH_3F molecules emerging from a quadrupole bent electrostatic guide [35]. A capillary was used to collimate and direct a thermal ($T = 298\,\mathrm{K}$) effusive beam of CH_3F towards the electrostatic guide, where only the slowest molecules were confined and guided around the bend whilst the fast-moving ones were lost and pumped away. Thus, the CH_3F molecules that collided with the Ca^+ ions were translationally cold but had a thermal distribution of rotational and vibrational states. Once the reaction occurred, the neutral product, CH_3, was lost but the ionic product, CaF^+, was trapped and sympathetically cooled by the Ca^+ Coulomb crystal such that the reaction could be monitored as a function of time by imaging the crystal with a CCD camera. The reaction was studied as a function of collision energy by varying the voltages applied to the quadrupole guide electrodes, and therefore varying the mean kinetic energy of the CH_3F reactant molecules between 2.0 and 3.7 K (compared to the ions' kinetic energy of 0.4 to 3.5 K depending on the crystal size). Despite the very low densities of CH_3F reaching the trap region (10^5 to $10^6\,\mathrm{cm}^{-3}$), the single ion sensitivity of the Coulomb crystal imaging set-up allowed the reaction to be investigated.

Recently, the reaction between Ca^+ ions within a Coulomb crystal and NO radicals has been reported [36]. The NO molecular radical is atypical since it is a stable gas and does not need to be produced *in situ* from a precursor. As such, it was introduced into the ion trap chamber through a leak valve in concentrations of 10^7 to $10^8\,\mathrm{cm}^{-3}$. The formation of the ionic reaction products NO^+ and CaO^+ was monitored with a CCD camera and probed by time-of-flight mass spectrometry as a function of the reaction time. The branching ratios and reaction rate constants for the charge exchange and the oxygen abstraction channels were determined at the average collision energy of 180 K. Additionally, the rate of both reactions was controlled by varying the quantum-state population of the Ca^+ ions in the Coulomb crystal. This was achieved by detuning the principal cooling laser to change the population of the $P_{1/2}$ state (the only reactive state populated) between 5% and 25%, which varied the rate of reaction by up to a factor of 5. Whilst the quantum-state population of the ions can be finely tuned, the control that has been achieved over the radical reactants is limited so far; NO has a thermal distribution of ro-vibrational states and velocities. In order to gain control over the quantum states and velocities of several selected radical species, various groups [36–38] are planning to combine an ion trap with a Stark decelerator [39]. Stark deceleration – the electric equivalent of Zeeman deceleration (described in the next section) – is a technique that employs pulsed electric fields to slow down and control the velocity of a beam of polar molecules in a quantum-state-selective manner [5]. Polar molecular radicals such as OH [40], NH [41], NO [42] and SD [43] have successfully been controlled using Stark decelerators and subsequently trapped electrostatically (OH [44] and NH [45]) or magnetically (OH [46]).

Ion traps have also been spatially overlapped with magneto-optical traps (MOT) to create hybrid traps for the study of ion–neutral interactions [47, 48]. Inside a MOT, the neutral species are confined by the quadrupolar magnetic field generated by a pair of coils in anti-Helmholtz configuration, and laser cooled [47]. Radical species such as Li [49], Na [50] and Rb [51] have been co-trapped with ions in hybrid traps. Rb atoms have been observed to react with a Ca^+ Coulomb crystal by charge exchange to form Rb^+ as well as by radiative association to form $CaRb^+$ [51]. The products were detected by resonant-excitation mass spectrometry of the Coulomb crystals. Subsequently, the kinetic energy of the Rb atoms was controlled and tuned (between 30 and 350 mK) by pushing the MOT cloud through the Coulomb crystal using controlled radiation pressure forces [52].

Although laser cooling of alkali-metal atoms and Stark deceleration of polar molecular radicals have proven to be effective methods to control the motion and quantum states of some radical species, not all radicals are suitable for laser cooling or can be manipulated with electric fields. In the following section, alternative methods for the control of radical beams are discussed.

1.2.2 Radical manipulation

Radicals are atomic or molecular species with one or more unpaired electron(s) in the outermost orbital that is responsible for their high reactivity. They are formed in very low densities under typical laboratory conditions, and they can be difficult to isolate due to (in most but not all cases) their lack of a charge and of an appreciable electric dipole moment. As a result of the presence of the unpaired electron(s), however, radicals always possess a magnetic dipole moment and therefore, in principle, they can all be manipulated using magnetic fields. In the presence of an external magnetic field, the energy levels of paramagnetic species split into high-field-seeking (HFS) and low-field-seeking (LFS) states according to the Zeeman effect (see section 2.1 for an illustration with respect to the H atom); the former states decrease whilst the latter increase in energy with increasing magnetic-field strength. Radicals in LFS states can be filtered out of an effusive source using magnetic guides, such as hexapole guides, where the LFS particles are transversely confined to the guide centre where the field minimises, whilst the HFS particles are lost [2]. If the guide is curved, only LFS radicals with a velocity below a certain threshold are filtered out, since the kinetic energy of the faster particles exceeds the depth of the confining potential so that they escape the guide by overshooting the bend [53]. Thus, a curved guide acts as a low-pass velocity filter, as well as a LFS-state filter. The guided radicals can be made internally cold, as well as translationally cold, and their flux can be increased by thermalising the sample with a buffer gas before guiding [54].

Whereas particles in effusive beams maintain the thermal internal energy and velocity distributions of the source, internally cold particles can be obtained from supersonic beams. Indeed, when the mean free path of the particles becomes smaller than the orifice from which they are escaping, particles collide frequently during the expansion and adiabatic cooling takes place. The decrease in translational (velocity spread), rotational and, to a lesser extent, vibrational energy is compensated by an increased kinetic energy of the beam in the laboratory frame [5]. There are two possible approaches to mitigating the high forward velocities of supersonic beams to study reactions at low collision energies. The first consists of crossing two supersonic beams at a very shallow angle such that the relative velocity of the cold species within each beam is minimised, as described below. The second approach involves removing the superfluous kinetic energy by decelerating the supersonic beam, and will be the focus of this thesis.

In order to reach collision energies corresponding to temperatures below 1 K, two supersonic beams need to be crossed at or very near to a 0° angle, practically meaning that they need to be merged into a single beam. This can be achieved by bending one beam into another (for cases where one of the species has no magnetic moment) employing a curved magnetic guide that, besides merging the beams so that they converge to the same axis, also provides state selectivity, given that only LFS particles are guided [55]. State selectivity of both beams can be attained by using two curved guides, such as a magnetic and an electric guide, to bend both beams so that they overlap [56]. The collision energy is controlled by varying the velocity of one of the beams and, in this way, collisions down to 8.7±0.8 mK have been recorded, allowing the observation of orbiting resonances in the Penning ionisation reaction of Ar and H_2 with metastable He [55]. Although the collision energy is normally limited by the velocity distribution of the hottest beam, which is around 0.1 K at best, lower energies can be obtained by using very short gas pulses generated by Even-Lavie valves [57]. When the flight time of the gas packet is considerably longer than the initial duration of the pulse, the position of the particles within it is correlated to their velocity: the fast particles travel in front and the slow ones lag behind. Selectively addressing a narrow portion of this velocity distribution in space results in lower collision energies and better resolution than would be possible otherwise [58]. Merged beams have been employed to investigate Penning ionisation reactions involving metastable helium [55] and neon [56], but the technique is also applicable to radicals, at least in principle. Despite this, studying ion–radical reactions with merged beams could prove to be challenging since the motion of the charged species would be affected by the magnetic fields needed to guide the radicals. With a view to improving the sensitivity of merged beam experiments, a magnetic synchrotron [59], analogous to the experimentally-verified electric version [60], has been proposed. In the synchrotron, particles are confined in packets within the storage ring and they interact with a tangential supersonic beam multiple times over numerous round trips, leading to an enhanced ability to discern between signal and noise.

An alternative approach to reaching low collision energies involves the reduction of the forward velocity of a beam of radicals after it has been supersonically expanded, which can be achieved using pulsed magnetic fields [61, 62]. Given that supersonic beams are collision-less environments, an increase in the potential energy of a particle is necessarily accompanied by a decrease in its kinetic energy and, therefore, a LFS particle travelling towards a region of high magnetic field gets decelerated. As soon as the particle travels past the region of maximum magnetic field, towards a low-magnetic-field region, its potential energy decreases again and the particle gets accelerated. If the magnetic field is switched off before the particle starts accelerating, the kinetic energy lost during the deceleration stage can be permanently removed. Following this principle, in a Zeeman decelerator (or coilgun) [63, 64] LFS particles in a supersonic beam are repeatedly obliged to travel through regions of positive magnetic-field gradient, whilst avoiding regions of negative magnetic-field gradient, so that their forward velocity is progressively diminished. The magnetic field is created by pulsing current through solenoid coils, with the beam travelling through their centre. Tens to hundreds of stages are required to bring the mean velocity of the particles to zero (depending on the mass of the radical and the magnetic moment of the specific quantum state to be decelerated), after which it is possible to confine the decelerated species within a magnetic trap [65]. However, only a small fraction of the LFS particles in the beam are typically successfully decelerated: those with a favourable combination of position and velocity that leads to a stable deceleration trajectory of the particle inside the decelerator. The region of phase space covering these favourable combinations is referred to as the phase-space acceptance of the decelerator; a large acceptance is desirable since high number densities of decelerated particles are experimentally advantageous. Various species have been addressed by Zeeman decelerators, including the radicals H [64], D [66], O_2 [67], CH_3 [68], Li [69], O [70] and the metastable species Ne [63], Ar [71], He_2 [72], He [73] and N [74]. Typical densities of decelerated particles achieved in conventional decelerators are around $10^8\,\mathrm{cm}^{-3}$ [62].

The periodic longitudinal (axial) deceleration experienced by the particles inside a Zeeman decelerator is accompanied by a periodic transverse (radial) focusing and de-focusing, as described in more detail in chapter 2. For particles travelling at low velocities (below $100\,\mathrm{ms}^{-1}$), the strong transverse focusing that they experience at the centre of each solenoid coil can become detrimental as it can deflect their trajectory and lead to particle loss. This transverse confinement issue is addressed in moving magnetic trap (or travelling wave) decelerators [71, 75] by confining and transporting particles within a three-dimensional travelling potential well. Instead of individual solenoid coils pulsed sequentially, a moving trap decelerator is composed of a sequence of spatially overlapped traps, each made of a pair of anti-Helmholtz coils [75]. The traps are activated sequentially with a temporal overlap such that the minimum magnetic field travels smoothly through the decelerator while transporting the trapped particles. By modulating the temporal overlap of the active traps, deceleration is achieved. The improved transverse confinement of the particles in this type of decelerator leads to larger phase-space acceptance and efficient trapping. Indeed, whilst standard Zeeman decelerators are more suitable for collision experiments, moving trap decelerators are optimal for trapping experiments [76]. Simultaneous deceleration of different species inside a moving trap decelerator (O_2 with metastable Ar [77]), followed by trapping within a permanent magnetic quadrupole trap (O_2 with Li [69]) has been demonstrated, opening the possibility of evaporative cooling or sympathetic cooling of molecules by co-decelerated laser-cooled atoms within a trap.

1.3 Aims

In order to study the reaction rate constants, temperature dependencies and product branching ratios of ion–radical reactions in a state-selective manner, the preparation of the reactants and the detection of the products need to be state selective. This is a challenging task but a necessary starting point is to create a source of cold state-selected radicals, which is the main focus of this thesis. The requirements for such a source are that it generates a pure beam of radicals (free of precursor molecules or other species) with control over the quantum state and velocity distribution of the particles, and with sufficient number densities such that an appreciable number of reactive collisions will occur on an appropriate experimental timescale. A highly sensitive detection technique is also needed to monitor product formation. The latter requirement will be satisfied by combining the radical source with a cryogenic linear Paul ion trap that is currently under construction. Sympathetically cooled molecular ions within Coulomb crystals will be the target for the beam of radicals, enabling the detailed study of reactions such as $H + CO_2^+$ (see section 1.1.1). The very low density with which radicals are typically formed under laboratory conditions can be partially compensated by the advantages of using Coulomb crystals as reaction targets, given the long interaction times and the very high sensitivity in the product detection that they afford (see section 1.2.1).

Generating a controllable and pure source of molecular free-radicals or open-shell atoms has been one of the primary barriers hindering the detailed study of radical processes in the laboratory. In this thesis, a velocity-selected radical beam source with a wide range of velocity tuneability and quantum-state selectivity is presented. The source is composed of a 12-stage Zeeman decelerator coupled with a newly-designed magnetic guide. The principles of Zeeman deceleration and the experimental apparatus are introduced in chapter 2. The first demonstration of the deceleration of metastable N atoms is presented in the same chapter (reported in [74]). The requirements that the radical source yield sufficient number densities and a pure beam of radicals are addressed in chapters 3 and 4, respectively. An evolutionary-algorithm approach is adopted to optimise the timed sequence of magnetic fields applied to the Zeeman decelerator in order to achieve maximum throughput of decelerated H radicals in chapter 3 (reported in [78]). A novel magnetic guide that purifies the decelerator output to generate a velocity-tuneable pure beam of velocity-selected H radicals is demonstrated in chapter 4 (reported in [79, 80]). Conclusions are drawn and implications and future work are discussed in chapter 5.

2 | Zeeman Deceleration

2.1 Principles of Zeeman deceleration

Zeeman deceleration is an experimental technique that employs magnetic fields to decrease the forward velocity of a supersonically expanded beam of paramagnetic particles, such as radicals or metastable electronically-excited species. The adiabatic supersonic expansion of a gas through a nozzle from a high pressure reservoir into a vacuum transforms translational, rotational and, to a lesser extent, vibrational energy into kinetic energy along the propagation direction of the beam. This process cools the internal and external degrees of freedom in the moving frame, at the expense of a high forward velocity in the laboratory frame [61]. In the type of decelerator used in this work, a tailored sequence of magnetic-field pulses is then applied to remove the excess lab-frame kinetic energy from the beam by making use of the Zeeman effect. The beam passes through a series of solenoid coils to which short pulses of high currents are applied in order to generate a large magnetic field inside the coils. The presence of unpaired electrons in paramagnetic species creates a magnetic dipole moment that couples to the applied magnetic field, lifting the degeneracy of the magnetic sub-levels. Those levels for which the energy increases with the applied field are known as low-field-seeking (LFS) states, whilst the states for which the energy decreases with field are known as high-field-seeking (HFS). Zeeman deceleration takes advantage of this interaction that particles in different magnetic sub-levels have with time-varying, spatially inhomogeneous magnetic fields to selectively address a subset of them: only paramagnetic species in selected LFS quantum states are decelerated [3, 61].

The Oxford Zeeman decelerator [81] is composed of a series of 12 solenoid coils that are rapidly pulsed with a current of 243 A, creating a maximum on-axis magnetic field of up to 1.8 T at the centre of each coil (see section 2.2.1). In the presence of an external magnetic field, the energy levels of atomic hydrogen are split: in ground-state hydrogen 1s ($^2S_{1/2}$), the interaction between electron spin ($S = 1/2$) and nuclear spin ($I = 1/2$) gives

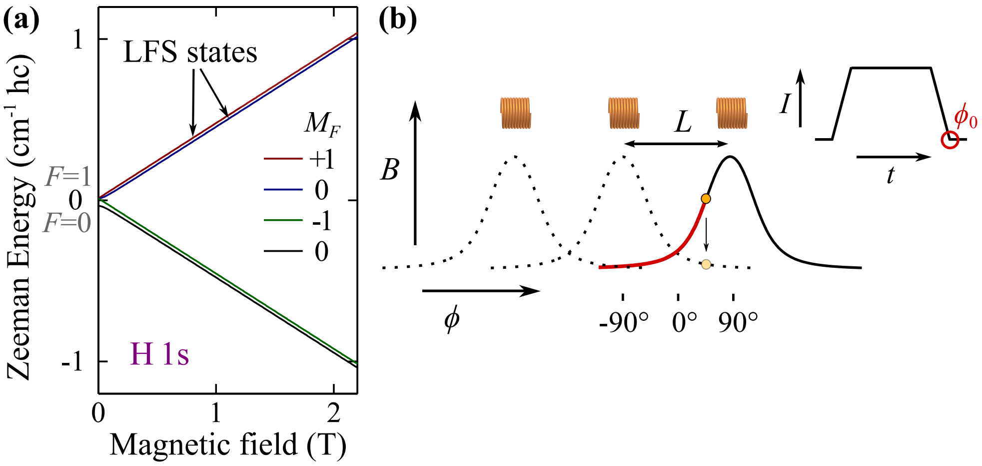

Figure 2.1: (a) Zeeman splitting for hydrogen $1^2\mathrm{S}_{1/2}$ atoms, with the low-field-seeking (LFS) states that are targeted during deceleration indicated. Specifically, these are the $M_F = 1$ and $M_F = 0$ states, where M_F is the projection of the total angular momentum F onto the local magnetic-field axis. (b) Schematic illustration of the magnetic-field profile as a function of the phase angle ϕ for three consecutive coils. ϕ indicates the position of a particle with respect to the centre of the coil (defined as $\phi = 90°$), whilst ϕ_0 denotes the position of the particle when the coil is switched off (i.e. the position of the particle when the current reaches zero – termed the effective phase angle – as shown in the inset).

rise to hyperfine splitting, yielding two LFS and two HFS states (Figure 2.1(a)). As the beam of paramagnetic particles travels towards the centre of each coil, particles experience a positive magnetic-field gradient and the potential energy of particles in LFS states increases at the expense of their kinetic energy (Figure 2.1(b)). These LFS particles can be permanently slowed down if the current (and therefore the magnetic field) is switched off before the particles reach the centre of the coil. The relative position of the particles with respect to the centre of the nearest coil is termed the phase angle ϕ, and is defined to be equal to 90° at the centre of a coil and 0° between two coils. The effective phase angle, ϕ_0, is used to indicate the position of the particles inside the coil at the time the magnetic field is switched off (Figure 2.1(b) inset).

The amount of kinetic energy removed per coil depends both on the magnetic moment of the quantum state and on ϕ_0. Specifically, each coil removes an amount of kinetic energy (ideally) equal to the Zeeman energy of the state for the specific magnetic field at the position of the particle when the coil is switched off. For example, for a particle in the

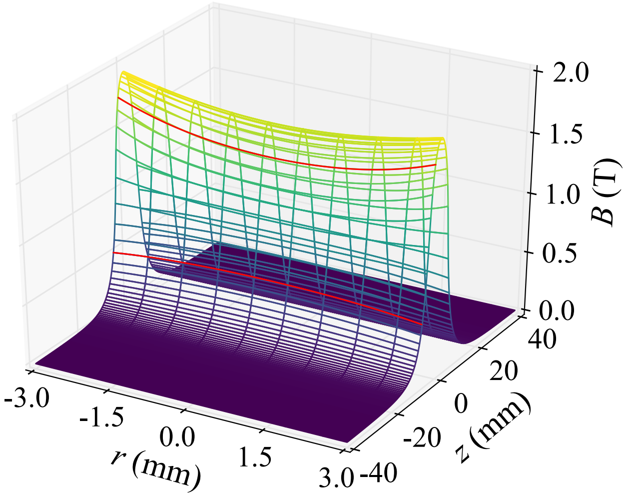

Figure 2.2: Magnetic field (in Tesla) generated by one solenoid coil as a function of longitudinal (z) and transverse (r) distance from the centre of the coil ($r = z = 0\,\mathrm{mm}$). The two red lines highlight the concave (de-focusing) transverse magnetic field turning convex (focusing) towards the coil centre.

$M_F = 1$ state that, at the time when the magnetic field is switched off, finds itself at the position $\phi_0 = 90°$, where the field is at its maximum value of $1.8\,\mathrm{T}$, the amount of kinetic energy removed is $0.84\,\mathrm{cm}^{-1}$hc (1.67×10^{-23} J) from Figure 2.1 (a). This idealised example assumes an instantaneous switching of the magnetic fields whilst, in reality, the current pulses through the coils have a finite rise and fall time, as described later in more detail (see section 2.2.1). Maximum deceleration is achieved with $\phi_0 = 90°$. A lower ϕ_0 value leads to less kinetic energy being removed, and therefore higher final velocities, whilst a value of $\phi_0 > 90°$ leads to re-acceleration of the particles. Thus a precisely timed switching sequence applied to the coils can slow a packet of quantum-state-selected H particles to an arbitrary final velocity, subject to the limitation of the number of coils available. The combination of initial positions and velocities of the particles that lead to stable deceleration trajectories inside the decelerator is referred to as the phase-space acceptance of the decelerator.

The number of decelerated particles transmitted by the decelerator as function of phase angle ϕ_0 is not constant; it is governed by the interplay between longitudinal and transverse focusing of the particles inside the decelerator. The magnetic field at the entrance of a coil is concave in the radial direction and therefore the particles are transversely de-focused at low ϕ_0 values (Figure 2.2). Increasing ϕ_0 towards 90° allows the particles to travel further inside the coil (before the magnetic field is switched off), where they experience transverse focusing due to the radial field becoming convex. The enhanced transverse focusing at large ϕ_0 values is counterbalanced by the longitudinal de-focusing, due to the loss of particles that travel past $\phi = 90°$ and get accelerated. Hence, the number of successfully decelerated particles is typically maximised at $\phi_0 \approx 50°$ [82].

Operating a short decelerator at $\phi_0 < 90°$, however, significantly reduces the amount of kinetic energy that can be removed from the beam and is therefore a counter-productive approach to maximising the number of decelerated particles transmitted. Instead, the phase-space acceptance of the decelerator can be maximised for deceleration at $\phi_0 = 90°$, such that the maximum amount of kinetic energy is removed from the beam. As first demonstrated with a Stark decelerator [83], the interplay between longitudinal and transverse motion can reduce the phase-space acceptance of the decelerator and thereby limit the efficiency of the deceleration process. Operation in, for example, the $s = 3$ configuration (where only every third stage is used for deceleration, with the other stages focusing the beam) effectively decouples the longitudinal and transverse dynamics, increasing the number of decelerated particles that can be obtained. There are, however, limitations associated with using modes of operation such as $s = 3$; with some coils facilitating transverse focusing in place of deceleration, far less kinetic energy can be removed compared with the standard mode of operation. Indeed, it has been previously proposed that the $s = 3$ mode of operation is not particularly suitable for Zeeman deceleration [84, 85].

Non-standard modes of operation have improved the transverse phase-stability of Zeeman deceleration, increasing the density of particles confined to the plane perpendicular to the molecular beam axis by using a mid-decelerator coil to focus the particles instead of to decelerate them [86]. This was achieved by reversing the current direction in the focusing coil to create a quadrupolar field that improved the transverse confinement of the particles, at the expense of one fewer deceleration stage. A new Zeeman decelerator has also recently been constructed where the decelerating solenoid coils are alternated with magnetic hexapoles that focus the beam radially [76]. The acceptance of the decelerator is optimised for scattering experiments through the decoupling of the longitudinal deceleration and the transverse focusing. An alternative approach is described in chapter 3, where evolutionary algorithms are employed to optimise the deceleration pulse sequences.

2.2 Experimental apparatus

The experimental apparatus consists of two differentially pumped high-vacuum chambers separated by a skimmer (Beam Dynamics, 2.0 mm orifice diameter). The source chamber ($p \approx 1 \times 10^{-5}$ mbar during experiments) contains a pulsed valve (General Valve Series 99, 0.8 mm orifice diameter, opening time of $\approx 230\,\mu$s FWHM) for supersonic expansion of the beam. The detection chamber ($p \approx 4 \times 10^{-8}$ mbar during experiments) contains the Zeeman decelerator and a time-of-flight mass spectrometer, for deceleration and detection of the particles respectively (Figure 2.3(a)). In order to reduce the initial kinetic energy of the beam, the pulsed valve body is cooled by a flow of gaseous nitrogen which has passed through a reservoir of liquid nitrogen, and its temperature is regulated with a PID-controlled heater element (CAL 9500P). The initial velocity of the beam is also decreased by seeding the precursor gas in a heavier noble gas.

Hydrogen atoms are generated with a velocity distribution peaking around $500\,\mathrm{m\,s^{-1}}$ from a mixture of NH_3 and Kr (1:9 ratio, 4 bar) which is supersonically expanded from the pulsed valve (kept at 238(3) K) into a quartz capillary (1 mm inner diameter, 15 mm length) attached to the face plate of the valve. Inside the capillary, NH_3 molecules are photolysed to ground-state $^2S_{1/2}$ H atoms by an ArF excimer laser (Neweks PSX-100, 193 nm, 1.5 mJ/pulse). The H atoms and other photo-fragments expand out of the capillary and travel through the skimmer and the decelerator, where a subset of the hydrogen atoms (in LFS states) is decelerated.

After passing through the decelerator, ground-state H atoms are photo-ionised by a $(2+1)$ resonance-enhanced multi-photon ionisation (REMPI) scheme via the 2s $^2S_{1/2}$ state, and the resulting ions are detected by a microchannel plate (MCP) detector (held at -1500 V)

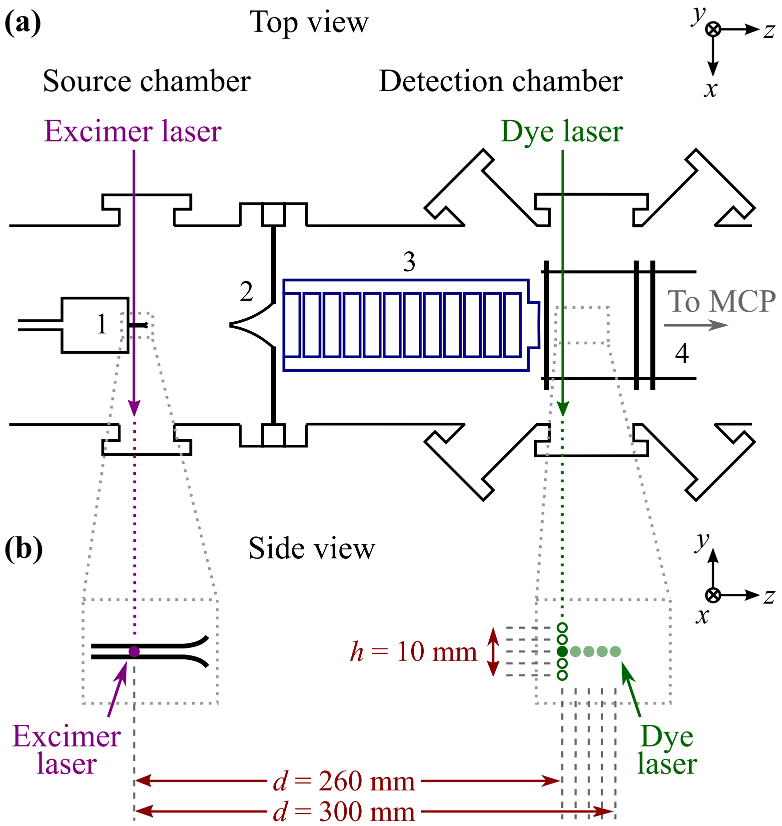

Figure 2.3: Schematic representation of the experimental apparatus (not to scale). (a) Top view: (1) pulsed valve, (2) skimmer, (3) Zeeman decelerator, (4) extraction plates. (b) Side view: only the magnified photolysis and photo-ionisation regions (delimited by dotted lines) are displayed. The measured ToF corresponds to the time taken by the H atoms to travel the distance d, which can be varied by translating the detection laser along the z axis (shaded green dots). The H beam profile can be probed by translating the detection laser along the y axis (empty green dots), away from the main beam axis position (full green dot). The range of horizontal detection positions used in section 3.3.1 ($d = 260$–300 mm), and the range of vertical detection positions used in section 4.2 ($h = 10$ mm around the beam axis located at $y = 0$ mm) are explicitly indicated.

in a Wiley-McLaren-type ion time-of-flight mass spectrometer. For the REMPI process, the 243.13 nm (1 mJ/pulse) frequency-doubled output of a dye laser (Spectra Physics PDL3, Coumarin 480 dye, 10 mJ/pulse) pumped by the third harmonic of an Nd:YAG laser (Quanta Ray, GCR 290, 10 Hz repetition rate) at 355 nm (130 mJ/pulse) is used.

The velocity distribution of the neutral H atoms is established by scanning the delay between the photolysis laser and the REMPI laser to collect traces that are referred to here as Time of Flight (ToF) profiles – with the ToF being the flight time of the H atoms between the axial positions where the beam is intersected by the two laser pulses (Figure 2.3). These ToF profiles are recorded both with and without the magnetic fields applied in alternating shots, achieved by switching the current-pulse sequence on and off. Therefore, whilst the repetition rate for the data acquisition is 10 Hz, the decelerator is usually operated at 5 Hz. Unless otherwise stated, the arrival time is scanned at 10 µs intervals, with 100 averages collected for each data point. The spatial distribution of hydrogen atoms leaving the decelerator is recorded by translating the REMPI laser beam with respect to the molecular beam axis by a micrometer-adjustable periscope. By moving the photo-ionisation laser horizontally, i.e. along the beam axis (Figure 2.3(b)), ToF traces at different photolysis-laser-to-detection-laser distances can be collected in order to experimentally probe the beam velocity (see section 3.3.1). In this way, the peak velocity of the H-atom beam when the magnetic fields are off is estimated to be $520 \pm 70\,\mathrm{ms}^{-1}$. By moving the photo-ionisation beam vertically, i.e. away from the beam axis, the H beam profile can be probed (see section 4.2).

2.2.1 Zeeman decelerator

The Oxford Zeeman decelerator is based on the design by the Merkt group [87], but with a novel adaptation that mounts the coils and cooling inside the vacuum (rather than outside) and allows interchange of single coil modules. Each of the 12 solenoid coils consists of Kapton-insulated copper wire (0.4 mm diameter) that is wound 70 times (over 4 layers) around a PEEK bobbin; the wound copper wire has an inner diameter of 7 mm, outer diameter of 10.6 mm and length of 8.5 mm. The exact dimensions of the solenoid have been established by matching the analytical solution for a current loop [88] to the magnetic field generated by a static current applied to the coils, measured using a Hall probe sensor (Honeywell, SS59ET) (see [81]). The wire is surrounded by thermally conductive epoxy resin and encased within an aluminium holder with notches to suppress eddy currents. The holder has an inner diameter of 6 mm and length of 10 mm.

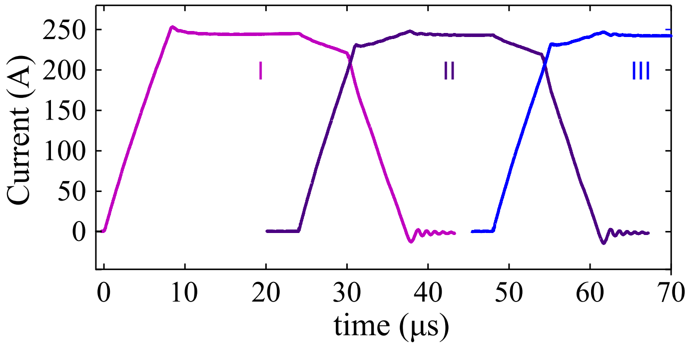

Figure 2.4: Experimental time profile of current pulses through coils I, II and III measured with a current probe. Each pulse has a rise and fall time of 7 µs, with a current plateau at 243 A in between. The 6 µs time overlap of adjacent coils prevents Majorana transitions to HFS states and causes mutual inductance effects resulting in the cusps present between the overlapped pulses. (Note that, due to the absence of mutual inductance, the rise time of coil I and the fall time of coil XII are 8 µs instead of 7 µs).

The 12 coils are mounted onto an aluminium base plate which is water cooled by a closed-cycle chiller system (ICS, Cygnus Tech, CY 031, water flow rate of 0.2 l/s; then replaced by Lauda MC1200, water flow rate of 0.5 l/s) to maintain coil temperatures below 330 K when pulsing 243 A of current at a 10 Hz repetition rate. The first coil is positioned 103.9 mm after the pulsed valve, and 54.9 mm after the skimmer; the distance between each coil is 0.7 mm, for a total length of the decelerator of 127.7 mm. The fast switching of the current pulses is achieved using an electronic circuitry based on insulated-gate bipolar transistors (IGBT), from the original design described in [87]. Current is supplied to the coils through two 12-pin power feedthroughs and the generated current pulses are monitored with a current probe (LeCroy CP500) (see Figure 2.4).

2.3 Numerical simulations

Three-dimensional numerical particle-trajectory simulations serve to both guide the selection of current switching sequences for the solenoid coils and to aid interpretation of the experimental ToF profiles. An arbitrary number of particles in specific quantum states can be simulated, with their initial positions and velocities randomly drawn from a normal distribution with parameters that match the experimental conditions. For the initial positions, the length and diameter of the capillary are used; the longitudinal and radial temperatures of the beam, as well as its mean longitudinal velocity, are determined by comparison with the experimental traces recorded with the magnetic fields off. Particles are assumed to be equally distributed amongst the four Zeeman sub-levels of H $1^2S_{1/2}$ and simulations are usually carried out for 50,000 particles in each state, unless otherwise specified. The geometry of the simulated set-up (the decelerator dimensions, current profiles, time-of-flight distances and laser beam dimensions) is as close as possible to the measured experimental parameters. Simulations are conducted with and without the magnetic fields present within the coils; when a current is applied to the coils, the resulting magnetic field (plotted in Figure 2.2) is derived from the analytical solution for a current loop [88].

In chapter 4, simulations are presented in which the post-decelerator beam is subsequently flown through the magnetic guide; the dimensions of all components (Halbach arrays and blades) are explicitly specified in the simulations. The Radia software package [89–91] is employed to model the magnetic fields arising from the Halbach arrays (see Figure 4.1(b)). Specifying the dimensions and magnetisation of the array, the software returns a three-dimensional grid of magnetic-field values as a function of position inside and outside the array. During the propagation of the particles, the magnetic field present at any arbitrary position is calculated by linear interpolation between the nearest points within the grid. The fields from all elements are summed vectorially to obtain the total magnetic field at any given position.

A velocity Verlet algorithm [92] integrates the equations of motion every 100 ns, with acceleration forces calculated using the particles' quantum state and the magnetic-field distribution. Particles are removed from the simulation when they encounter a surface (such as a skimmer, decelerator coil, Halbach array or blade) or reach the detection region, defined as the volume of space illuminated by the detection laser. The detection laser is aligned with the x axis (perpendicular to the atomic beam, which propagates along the z axis) and is modelled as a Gaussian in the z and y dimensions, with a standard deviation of 0.8 mm and 0.4 mm respectively. The simulated signal is given by the number of trajectories passing through the laser detection volume in a given time interval, and is used to generate ToF profiles for comparison with the experimental traces. For the optimisation of the guide geometry in chapter 4, the detection region at the end of the magnetic guide is defined as a circle with a 1.3 mm radius to match the acceptance of a linear Paul ion trap, for future ion–radical collision studies.

2.3.1 Generation of standard pulse sequences

The conventional, or standard, current-pulse sequence applied to the decelerator coils is intuitive: it consists of pulses of monotonically increasing duration applied to each coil in sequence (Figure 2.5(a)). This arises because standard switching sequences are calculated from the trajectory simulations (see section 2.3) considering the passage of a single "synchronous" hydrogen atom in the $M_F = 1$ state through the apparatus. The idealised synchronous particle ($x = y = 0\,\mathrm{mm}$ and $v_x = v_y = 0\,\mathrm{ms}^{-1}$) travels along the axis of the decelerator. When the synchronous particle reaches a desired axial position within each coil, the magnetic field within the coil is switched off. This is repeated for all twelve coils, and in this way a complete switching sequence is constructed. As the synchronous particle travels through each coil, it gets increasingly decelerated and therefore it takes longer to reach the switch-off position in the next coil, thus giving rise to the sequential increase in successive pulse durations.

As can be seen in Figure 2.4, the current pulses applied to the coils in the apparatus exhibit non-zero rise and fall times of 7 µs. This is because the magnetic-field strength changes at a finite rate, dictated by the time constant of the inductor-resistor circuit formed by the coil and its switching electronics. The switch-on time of each coil is chosen so that the rising field of a coil is overlapped in time with the decaying field of the previous coil; an overlap time of 6 µs is adopted here. This temporal overlap of adjacent pulses maintains a non-zero quantization magnetic field, thus avoiding Majorana spin-flip transitions to HFS states [66]. Maximum deceleration is achieved at a phase angle of $\phi_0 = 90°$, for which a synchronous particle in the $M_F = 1$ Zeeman sub-level with initial velocity of $500\,\mathrm{m\,s}^{-1}$

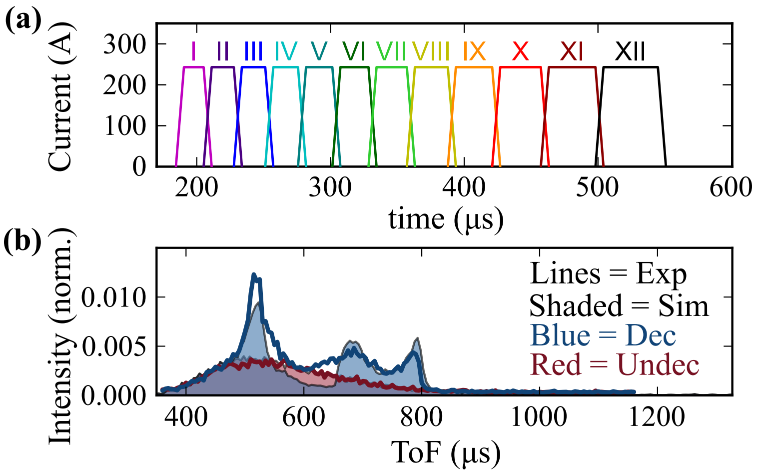

Figure 2.5: (a) Schematic illustration of a standard pulse sequence calculated from trajectory simulations: the time profile of the current passed through each of the coils (labelled with roman numerals) is shown. The standard sequence is generated using a phase angle $\phi_0 = 90°$, which decelerates the synchronous particle from $v_z^{\mathrm{initial}} = 500\,\mathrm{m\,s^{-1}}$ to $v_z^{\mathrm{final}} = 205\,\mathrm{m\,s^{-1}}$. (b) Time-of-flight (ToF) traces obtained when applying the sequence in panel (a) to the decelerator and detecting the H atoms at the detection distance $d = 260\,\mathrm{mm}$ (see Figure 2.3(b)). Solid lines indicate the experimental traces and shaded areas show the simulated traces. For each ToF profile, both the "undecelerated" (magnetic fields off, red colour) and "decelerated" (magnetic fields on, blue colour) profiles are shown.

is slowed to a final velocity of $205\,\mathrm{m\,s^{-1}}$ after 12 coils. When a standard pulse sequence is implemented experimentally, a subset of particles with positions and velocities similar to the idealised synchronous particle lie in a phase-stable region and are also decelerated, forming a compact bunch that travels through the decelerator. The size of this phase-stable region depends on the choice of ϕ_0: the closer ϕ_0 is to 90°, the greater the kinetic energy that is removed at each stage but, at the same time, the smaller the size of the stable phase-space volume, which means that fewer particles are decelerated (see section 2.1).

Figure 2.5(b) shows the experimental and simulated ToF traces obtained when the standard sequence in panel (a) is applied to the decelerator coils. ToF traces recorded with no deceleration (i.e. with no current applied to the coils, referred to here as "undecelerated" traces) are shown in red, with the experimental trace shown as a solid red line. The broad, featureless profile of the undecelerated beam is in excellent agreement with the simulated ToF profile (shaded red trace), confirming that the properties of the initial beam and the decelerator geometry are well described in the simulations. The mean longitudinal velocity and temperature of the undecelerated beam are estimated from the simulations to be $480\,\mathrm{m\,s^{-1}}$ and 2.0 K, respectively; the radial temperature is estimated to be 0.1 K. The "decelerated" ToF traces (i.e. where current is applied to the coils, blue profiles) show a characteristic multi-peak profile. The first and most intense peak (with the earliest arrival time) corresponds to the part of the beam that gets focused as it travels through the decelerator, without being significantly decelerated or accelerated. The final peak (with the latest arrival time) corresponds to the part of the beam that gets decelerated to the target velocity. The features between the first and final peaks in the decelerated ToF traces arise from particles that are partially decelerated by the magnetic fields. The multiple peaks and features of the experimental decelerated ToF profile (solid blue line) are well described by the simulations (shaded blue profile); any minor discrepancies between the simulated and experimental ToF profiles are attributed to differences between the simulated ideal gas pulse and the experimental beam.

2.4 Deceleration of metastable nitrogen atoms

In this section, the first demonstration of the Zeeman deceleration of metastable N atoms is presented. In a proof-of-principle experiment, a velocity decrease from 460 to $410\,\mathrm{ms^{-1}}$ is achieved. Excited-state atoms and molecules are referred to as being "metastable" when their relaxation to the electronic ground state is formally forbidden by electric-dipole selection rules, thus resulting in natural lifetimes that exceed the nanosecond range by several orders of magnitude. Atomic nitrogen has an exceptional energy-level structure, as it possesses a metastable $^2\mathrm{D}_J$ state that lies just 2.4 eV above the electronic ground state [93]. The lifetime of the $J = 5/2$ ($J = 3/2$) spin-orbit component is estimated to be 42 hours (17 hours) [94], since decay to the $^4\mathrm{S}_{3/2}$ ground state is only possible via electric quadrupole (magnetic dipole) transitions.

The hyperfine interaction and coupling between the electron spin ($S = 1/2$) and the

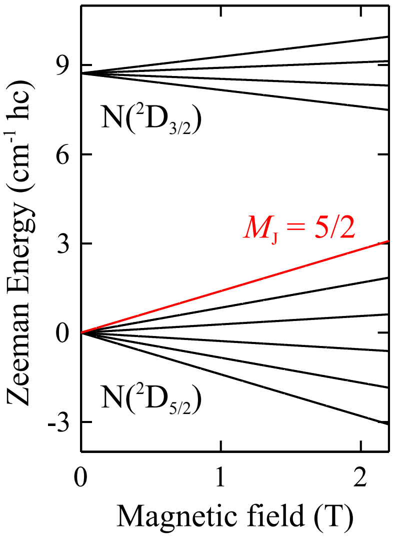

Figure 2.6: Zeeman diagram for N(^{2}D$_J$) with $J = 3/2$ and $5/2$. The $M_J = 5/2$ energy sub-level used for the calculation of the Zeeman deceleration pulse sequence is highlighted in red.

nuclear spin ($I = 0$) in N(^{2}D) is negligible [95]. Therefore, the Zeeman shift due to the interaction of the electron spin with an external magnetic field can be expressed by $\Delta E_\mathrm{Z} = M_J g_J \mu_\mathrm{B} B$, where M_J is the projection of the total angular momentum J onto the local magnetic-field axis, μ_B is the Bohr magneton and B denotes the magnetic field. The Landé factor, g_J, amounts to 1.20 for the ^{2}D$_{5/2}$ state and 0.80 for the ^{2}D$_{3/2}$ state [96]. The Zeeman diagram of the ^{2}D state (Figure 2.6) illustrates that the Zeeman splitting is quite large, especially for the $J = 5/2$, $M_J = 5/2$ state, for which the energy shift at $B = 2\,\mathrm{T}$ amounts to about 2.8 cm^{-1}.

2.4.1 Methods

The apparatus used for these experiments is the same as the one described in section 2.2 for the generation and deceleration of H atoms, with the following modifications. Metastable ^{2}D$_{5/2}$ nitrogen atoms are produced from N$_2$ seeded in Ar (1:6 ratio, $p_0 = 5.5\,\mathrm{bar}$) expanded through the pulsed valve (kept at $T = 133(3)\,\mathrm{K}$) and dissociated by electron impact (perpendicular to the gas beam) using a home-built, pulsed electron gun. The experiments described here are carried out at an average kinetic energy of 100 eV

($I_{\text{fil}} = 3.92\,\text{A}$) using electron-pulse lengths of either $20\,\mu\text{s}$ or $50\,\mu\text{s}$. The capillary is not present in this version of the setup and the valve-to-decelerator distance is decreased by 13 mm compared to the hydrogen experiments. The distance from the valve to the detection laser, d, is fixed at 255 mm (see Figure 2.3).

The detection of metastable N in the $^2\text{D}_{5/2}$ state is achieved via a $(2+1)$ REMPI scheme using two photons at 268.97 nm (2.5 mJ/pulse) for resonant excitation to the intermediate 3p $^2\text{S}_{1/2}$ state, and a third photon of the same energy for photo-ionisation. Laser radiation at this wavelength is generated by frequency-doubling the output of the dye laser operated with Coumarin 540A. A continuous voltage is applied to the extraction plates to prevent the detection of spurious ions produced by the electron-impact excitation process. The voltage of the MCP detector is gated between $-1530\,\text{V}$ and $-1650\,\text{V}$ to allow for a better discrimination between the N^+ ion signal and noise.

A deceleration pulse sequence is calculated using trajectory simulations to selectively decrease the longitudinal velocity of a synchronous particle in the $^2\text{D}_{5/2}, M_J = 5/2$ sub-level moving initially at $460\,\text{ms}^{-1}$ down to $410\,\text{ms}^{-1}$. Alternatively, a second pulse sequence is used to decelerate the synchronous particle from $400\,\text{ms}^{-1}$ down to $340\,\text{ms}^{-1}$. Rather than calculating the trajectory of the synchronous particle for a constant phase angle, in this section the deceleration experienced by the synchronous particle is kept constant in every deceleration stage. Due to the finite rise and fall times (of $7\,\mu\text{s}$) of the current through the coils, deceleration at a constant phase angle does not remove an equal amount of kinetic energy from the synchronous particle at every stage. Since the magnetic field must be off when the particle reaches $\phi_0 = 90°$ to prevent acceleration, the current must be switched off $7\,\mu\text{s}$ before the particle reaches the middle of each coil (see Figure 2.1(b), inset). The position of the particle (and therefore the magnetic field it experiences) when the current is switched off varies based on the velocity of the particle: a fast particle will travel further in $7\,\mu\text{s}$ and therefore the current must be switched off when the particle is further away from the coil centre compared to a slower particle. As the deceleration proceeds, the synchronous particle is allowed to travel further into the coil before the current is switched off and so each consecutive coil removes slightly more kinetic energy than the previous one. Instead of a constant phase angle, a variable parameter $\kappa_0 = (z - z_0)/d$ is used here to maintain a constant (and maximal) deceleration of the synchronous particle at every stage. This dimensionless parameter relates the on-axis particle position relative to the coil centre, $z - z_0$, to the centre-to-centre coil distance, d. The implementation of this "adaptive phase angle" is not detailed here but is described in [84].

In the three-dimensional particle-trajectory simulations, 1.5 million particles are simulated in each of the six Zeeman sub-levels of the $^2\mathrm{D}_{5/2}$ manifold; the particle distribution among the sub-levels is assumed to be equal. In the results described below, all simulation input parameters are the same, except for the settings that are changed in between experiments: (a) the pulse duration for electron impact, which is varied between 20 µs and 50 µs and (b) the Zeeman-deceleration pulse sequences for deceleration from an initial velocity of 460 ms^{-1} or 400 ms^{-1}. The mean longitudinal velocity and temperature of the undecelerated beam are estimated from the simulations to be 470 $\mathrm{m\,s}^{-1}$ and 1.65 K respectively; the radial temperature is estimated to be 0.1 K.

2.4.2 Results

The top row in Figure 2.7 shows experimental ToF traces for Zeeman-deceleration experiments carried out using two different pulse durations for electron-impact excitation, 50 µs (panels (a) and (c)) and 20 µs (panel (b)), and using two different deceleration sequences, from 460 to 410 ms^{-1} (panels (a) and (b)) or from 400 to 340 ms^{-1} (panel (c)). Here, the time of flight is specified as the delay between the rising edge of the electron-impact-excitation pulse and the detection-laser pulse. Below the experimental ToF profiles, the corresponding results from numerical particle-trajectory simulations are displayed (panels (d)–(f)). To better visualise the change in the relative signal intensity for both the experimental and simulated traces, the difference in ToF intensity between decelerated (decelerator on) and undecelerated (decelerator off) profiles is shown in panels (g)–(i).

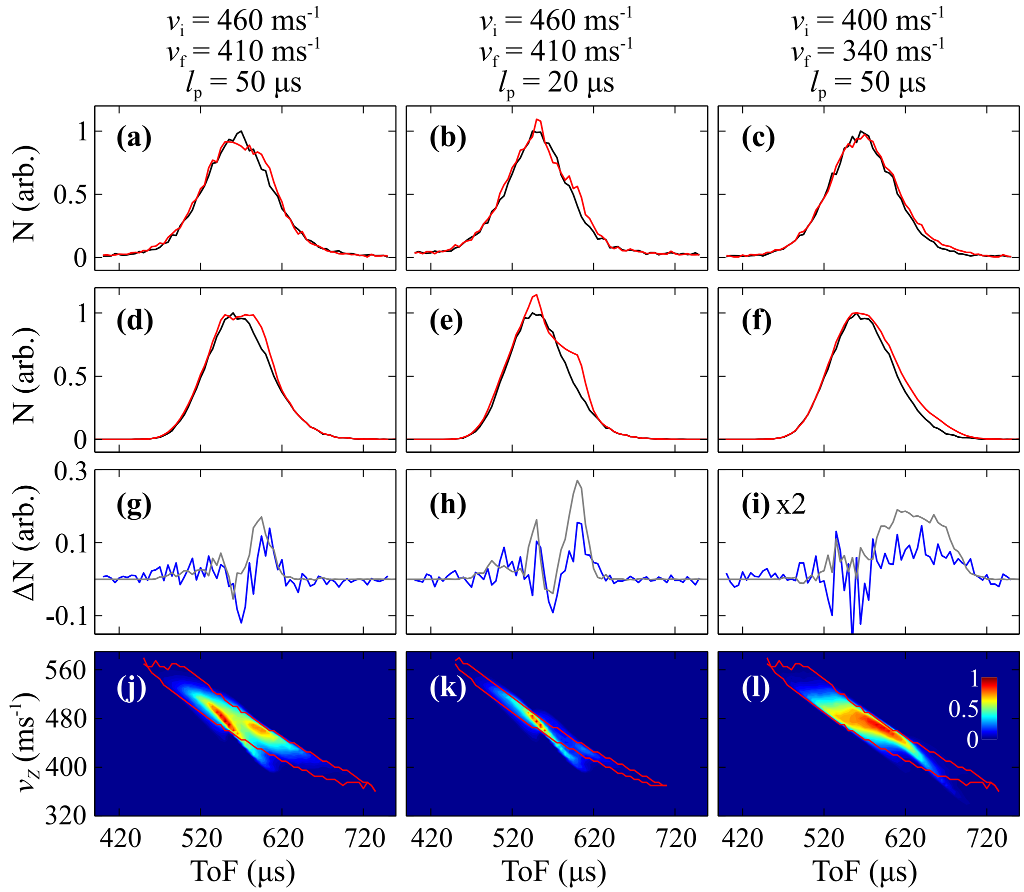

Figure 2.7: Experimental (a)–(c) and simulated (d)–(f) ToF traces for the deceleration of nitrogen atoms in the $^2D_{5/2}$ state using different deceleration sequences and electron-pulse durations. Above each column, the initial (v_i) and final (v_f) velocity of the deceleration sequence used, as well as the pulse length (l_p) for the election-impact excitation employed, are indicated. The undecelerated (magnetic fields off) and decelerated (magnetic fields on) profiles are shown in black and red, respectively. The experimental data is averaged over 15 measurements for 50 μs electron-gun-pulse durations and over 25 measurements for 20 μs electron-gun-pulse durations, with 100 averages per data point for each individual measurement. The difference between decelerated and undecelerated ToF traces for the experimental (blue) and simulated (grey) data is shown in panels (g)–(i) (note that (i) is magnified by a factor of two for better visibility). The simulated final longitudinal velocity (v_z) as a function of ToF is shown in panels (j)–(l); the red curve indicates the convex hull of the velocity distribution obtained without deceleration.

When particles with a longitudinal velocity of 460 ms^{-1} are selected for Zeeman deceleration (Figure 2.7(a) and (b)), the signal intensity of the decelerated beam (red) is slightly increased at ToF ≈ 600 μs compared to the undecelerated beam (black). The decelerated trace displays a shoulder, representing an underlying peak of decelerated particles, that is not present in the undecelerated trace. This change in signal intensity between decelerated and undecelerated beams as a function of flight time is reproduced by the particle-trajectory simulations (Figure 2.7(d) and (e)) and is most evident in the difference traces (Figure 2.7(g) and (h)), where the positions of the experimental signal maxima and minima (blue) are well-reproduced by the simulations (grey). There is also qualitative agreement between experimental and simulated traces regarding the relative magnitude of the maxima and minima. The decelerated shoulder is more pronounced using electron-impact-excitation pulses of 20 μs, compared to 50 μs. The longer electron-gun pulses produce a larger number of metastable N atoms, leading to better S/N ratio but also a broader spread of arrival times of the atoms, which obscures the decelerated signal more than a shorter pulse. In the case of deceleration from 400 to 340 $\mathrm{m\,s^{-1}}$ (Figure 2.7(c), (f) and (i)), the effect of the deceleration is less immediately apparent, as there is only a very small relative signal increase in the decelerated trace which is spread over the falling tail of the ToF profile (panel (c)). However, this small effect is predicted by the trajectory simulations (panel (f)), and there is qualitative agreement between experiments and simulations regarding the shape and timing of the decelerated signal (panel (i)).

The simulated ToF profiles, most notably in Figure 2.7(e), overestimate the signal enhancement stemming from the decelerated particles. A similar discrepancy has been previously observed in the deceleration of metastable He atoms [73] and is attributed to additional losses during or after the deceleration process that are not accounted for in the simulation [84]. Additionally, due to low REMPI signal intensities, data acquisition over several hours was required to improve the signal-to-noise ratio. Fluctuations in the experimental conditions, such as the valve backing pressure and the characteristics of the electron gun, may have contributed to the observed discrepancy.

Panels (j)–(l) in Figure 2.7 illustrate the final longitudinal velocity distribution of the simulated particles as a function of ToF for the three different experimental settings; the red curves indicate the boundaries of the velocity distributions obtained in the absence of the magnetic fields. The deceleration induces a depletion of particles travelling at the velocity selected for deceleration (460 $\mathrm{m\,s^{-1}}$ in (j) and (k), 400 $\mathrm{m\,s^{-1}}$ in (i)) and the appearance of particles outside the complex hull travelling at the expected final velocity (410 $\mathrm{m\,s^{-1}}$ in (j) and (k), 340 $\mathrm{m\,s^{-1}}$ in (i)).

There are several reasons for the absence of a bunch of decelerated particles that is clearly separate from the rest of the beam. Firstly, the higher atomic mass of N (compared to H) leads to a lower magnetic-moment-to-mass ratio, so that less quantum-state selectivity (transverse focusing and defocusing) and less deceleration are achieved with the same number of deceleration stages. Despite the relatively low initial mean supersonic-beam velocity of 470 $\mathrm{m\,s^{-1}}$ – attained through cooling of the pulsed valve body and seeding in Ar gas – deceleration from 460 to 410 $\mathrm{m\,s^{-1}}$ only amounts to a 20 % decrease in kinetic energy. Deceleration from 400 to 340 $\mathrm{m\,s^{-1}}$ results in a higher kinetic energy removal of almost 30 %. However, fewer particles within the initial velocity distribution (centred at 470 $\mathrm{m\,s^{-1}}$) are selected by the deceleration pulse sequence and, therefore, particle deceleration is even more difficult to observe under these conditions. Secondly, the short length of the decelerator limits the spatial separation of the particles according to their velocity, resulting in very little temporal delay between the arrival of the decelerated and undecelerated particles. Another contributing factor is the REMPI detection of metastable N atoms, which allows for the selective detection of the $^2\mathrm{D}_{5/2}$ spin-orbit component but does not distinguish between its different M_J states.

Even though both deceleration pulse sequences are specifically tailored to address particles belonging to the $M_J = 5/2$ state, the individual Zeeman sub-levels of the $^2\mathrm{D}_{5/2}$ manifold are not well-separated during the deceleration process. Hence, all six sub-levels contribute to the observed ToF profiles, as illustrated in Figure 2.8 where the initial and final velocity of the simulated particles are shown after deceleration from 460 to 410 $\mathrm{m\,s^{-1}}$ (same input parameters as in Figure 2.7(e), (h) and (k)). Particles in all Zeeman sub-levels of the $^2\mathrm{D}_{5/2}$ state are present and their velocity is decreased (or increased in the case of HFS states) according to their different Zeeman interaction with a magnetic field (Figure 2.6). For (high) low-field-seeking quantum states the final velocities are (highest) lowest for the largest absolute values of M_J. Hence, the particles are decelerated and accelerated to various different velocities, leading to the complex ToF profile in Figure 2.7(e).

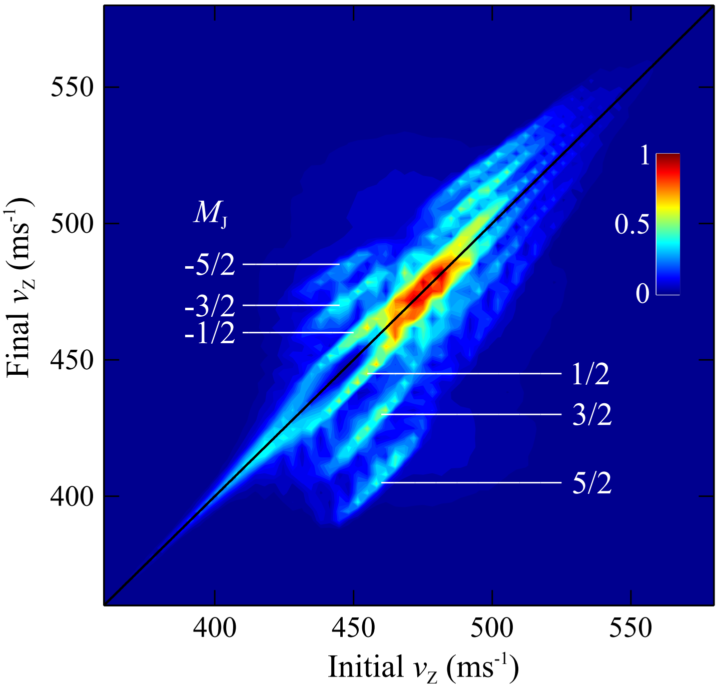

Figure 2.8: Simulated final longitudinal velocity as a function of initial longitudinal velocity after deceleration of N(^{2}D$_{5/2}$) from 460 ms^{-1} to 410 ms^{-1} with an electron-impact-excitation-pulse length of 20 µs (same input as panels (e), (h) and (k) in Figure 2.7). Only particles inside the volume of the detection laser are shown. A branching of the velocities corresponding to quantum states with different M_J values is clearly visible. The black diagonal line corresponds to the velocities that are unaffected by the deceleration process.

2.4.3 Conclusion

Despite the experimental challenges, the results presented here clearly illustrate that Zeeman deceleration of N(^{2}D$_{5/2}$) can be achieved. The use of a decelerator with a larger number of deceleration stages would greatly improve the experimental conditions for N-atom deceleration, requiring, for example, less cooling of the pulsed valve or less dilution in a heavy carrier gas. A longer decelerator would also improve quantum-state selectivity and increase the temporal separation between the decelerated and undecelerated parts of the beam. Zeeman deceleration from 460 ms^{-1} to a standstill could be achieved within 50 deceleration stages. On the other hand, optimising the deceleration pulse sequences to increase the number of decelerated particles – or even to improve the decelerative ability of the current decelerator – is also a possibility.

3 | Evolutionary algorithms optimisation

3.1 Principles of evolutionary algorithms

Conventionally, the deceleration switching sequence – that is, the timing of the switching on and off of the magnetic field inside each solenoid coil composing the decelerator – is determined by considering the trajectory of an idealised synchronous particle, as described in section 2.3.1. The trajectory of this synchronous particle under the influence of the magnetic fields is calculated, and the switching time of each solenoid coil is determined by when the particle reaches a selected axial position within that coil relative to the coil centre. This position is normally expressed as a phase angle and usually such devices are operated at constant phase angle (and constant maximum current) throughout the deceleration sequence. Changing the selected phase angle enables a range of different final velocities of the decelerated species to be attained, albeit with variable transmission efficiency, as discussed in section 2.1. The switching sequences produced by this standard approach are capable of decelerating other particles close to the synchronous particle in phase space, but they are not necessarily the optimal sequences; there are limitations to using standard switching sequences when seeking to control the properties of the resulting beam.

The use of an adaptive phase angle (varying through the sequence) to keep the deceleration per stage constant (as introduced in section 2.4.1) has been investigated, and it was demonstrated to enhance the decelerator acceptance when operating the device in acceleration mode, but not to give a significant improvement over the constant-phase-angle operation mode when decelerating [82, 84]. Given that the individual solenoid-pulse timings can be varied arbitrarily and independently in a given apparatus, more complex timing sequences could beneficially adapt the characteristics of the output beam for specific applications. In order to fully optimise the switching sequences to achieve the maximum number of decelerated particles, one must consider more than just the synchronous particle. A suitable approach to investigate the full optimisation of the sequence, subject to certain practical constraints, is to use an evolutionary strategy.

Evolutionary strategies are a class of numerical optimisation techniques, drawing inspiration from Darwin's evolutionary theory of natural (genetic) selection. Since genetic algorithms were popularised by Holland [97] as a way of solving computationally intractable problems, genetic and evolutionary algorithms have been implemented in a wide range of disciplines spanning the mathematical, physical and medical sciences. An abstract adaptation of biological genetic processes – selection, recombination, evolution and mutation – is employed to modify a set of parameters that influence a measurable outcome, and to locate the optimal solution for that set of parameters required to achieve the closest match to the desired outcome. From an initial set of parameters, an evolutionary algorithm carries out a process of fitness-based ranking and selection, followed by recombination to produce a successor population of parameter sets: the next generation. During recombination, the best performing parent "chromosomes" are preferentially selected and their "genetic material" is recombined and passed on to the successor population. As this process is iterated, successive generations are formed and the average fitness of the potential solutions typically improves until convergence is achieved (or some stopping criterion is reached). In this way, an evolutionary algorithm "evolves" the optimal solution to a given problem. Here, a covariance matrix adaptation evolutionary strategy (CMA-ES) is the evolutionary algorithm adopted [98].

Applying this approach to Zeeman deceleration, each trial solution comprises a set of values for the parameters that define the switching sequence of the coils (i.e. pulse durations and timings), and the fitness of a given sequence can be defined as the fraction of particles that exit the decelerator with a longitudinal velocity within, for example, $\pm 10\,\mathrm{m\,s^{-1}}$ of the target velocity. To achieve optimisation, looking at a single synchronous particle is not sufficient; the trajectories of the full set of particles representing the initial distribution in phase space through the decelerator apparatus need to be considered.

For a constant current of 243 A applied to the solenoid coils, optimisation of the pulse durations represents a 12-dimensional problem for the 12-coil apparatus in this thesis. While iterative steepest-decent methods, such as the Levenberg-Marquardt algorithm [99–101], are useful when the initial parameter set is close to optimal, stochastic methods are better suited at locating the global minimum of high-dimensional systems. Furthermore, evolutionary strategies do not require any information about the gradient of the target function at a given set of parameters, meaning that the function itself need not be continuous or differentiable. Evolutionary strategies are thus broadly applicable and typically cost-effective to implement. While evolutionary algorithms have already been shown to increase the number of decelerated particles transmitted in both Stark [102] and Zeeman [82] deceleration, their full potential was perhaps not fully exploited in these earlier studies. In Zeeman deceleration, a CMA-ES approach was adopted to maximise the number of decelerated particles whilst minimising their spread in velocity and position. An increase in the intensity of decelerated D atoms was reported and attributed to improved phase-space acceptance and optimal balance between transverse focusing and defocusing throughout the deceleration process. However, the evolutionary-strategy-optimised sequences were not found to be a significant improvement over the conventional constant-phase-angle sequences with phase angles in the range 45°–55° [82].

This work builds on previous studies and utilises the CMA-ES evolutionary algorithm to fully optimise the operation of a Zeeman decelerator for collision studies. This chapter describes how evolutionary algorithms can generate deceleration sequences that not only achieve more efficient deceleration than a comparable standard sequence – yielding a higher number of decelerated particles within a given narrow velocity range – but also enable significantly lower final velocities to be achieved for the decelerator used in this work than would be possible with a standard sequence. It demonstrates how evolutionary algorithms can be employed to enhance the ability to manipulate the resulting decelerated beam, exercising far more control over the beam than is possible through phase-angle manipulation alone. Finally, it investigates the deceleration mechanism employed by the standard and optimised sequences, identifying the key differences between the two approaches and how the optimised sequences yield such significantly improved results.

3.2 Implementation

3.2.1 Evolutionary optimisation of pulse sequences

Evolutionary algorithms seek to optimise a set of parameters with respect to a fitness function: the “fitness” of each set of parameters (which give rise to a deceleration sequence) in a population is computed, after which a new generation of sequences is created based on the results from the previous generation. The parameters that are optimised here are the durations of the pulses of current applied to each of the twelve coils. A full switching sequence is generated from a set of durations by applying the necessary constraints. Each pulse must overlap with the previous one to maintain a non-zero quantization magnetic field, as described in section 2.3.1 for the standard sequences (the same overlap time of 6 µs is adopted). Additionally, for the apparatus employed here, each pulse must not exceed 60 µs in duration. The upper limit of the pulse duration (60 µs) arises from the limited cooling efficiency of the solenoid coils, given the large amount of heat that needs to be dissipated when repeatedly pulsing such high currents (243 A) at 10 Hz. The lower limit of the pulse duration is 14 µs, given by the 7 µs rise and fall times of the current pulses. The total number of particles arriving within $\pm 10\,\mathrm{m\,s^{-1}}$ of the target velocity is optimised, with the fitness of each sequence evaluated by simulating the trajectories of 50,000 particles in the $M_F = 1$ Zeeman sub-level and recording the number that reach the detection region with a velocity within the target range.

A CMA-ES approach is used to optimise the switching sequence so as to maximise the number of decelerated particles reaching the laser detection volume. A standard pulse sequence (the sequence in Figure 2.5(a), yielding a final velocity $v_z^{\mathrm{final}} = 205\,\mathrm{m\,s^{-1}}$) is taken as the starting point to create the first generation of sequences with a slightly lower target final velocity $v_z^{\mathrm{final}} = 200\,\mathrm{m\,s^{-1}}$ in the evolutionary-algorithm approach (Figure 3.1(a)). A

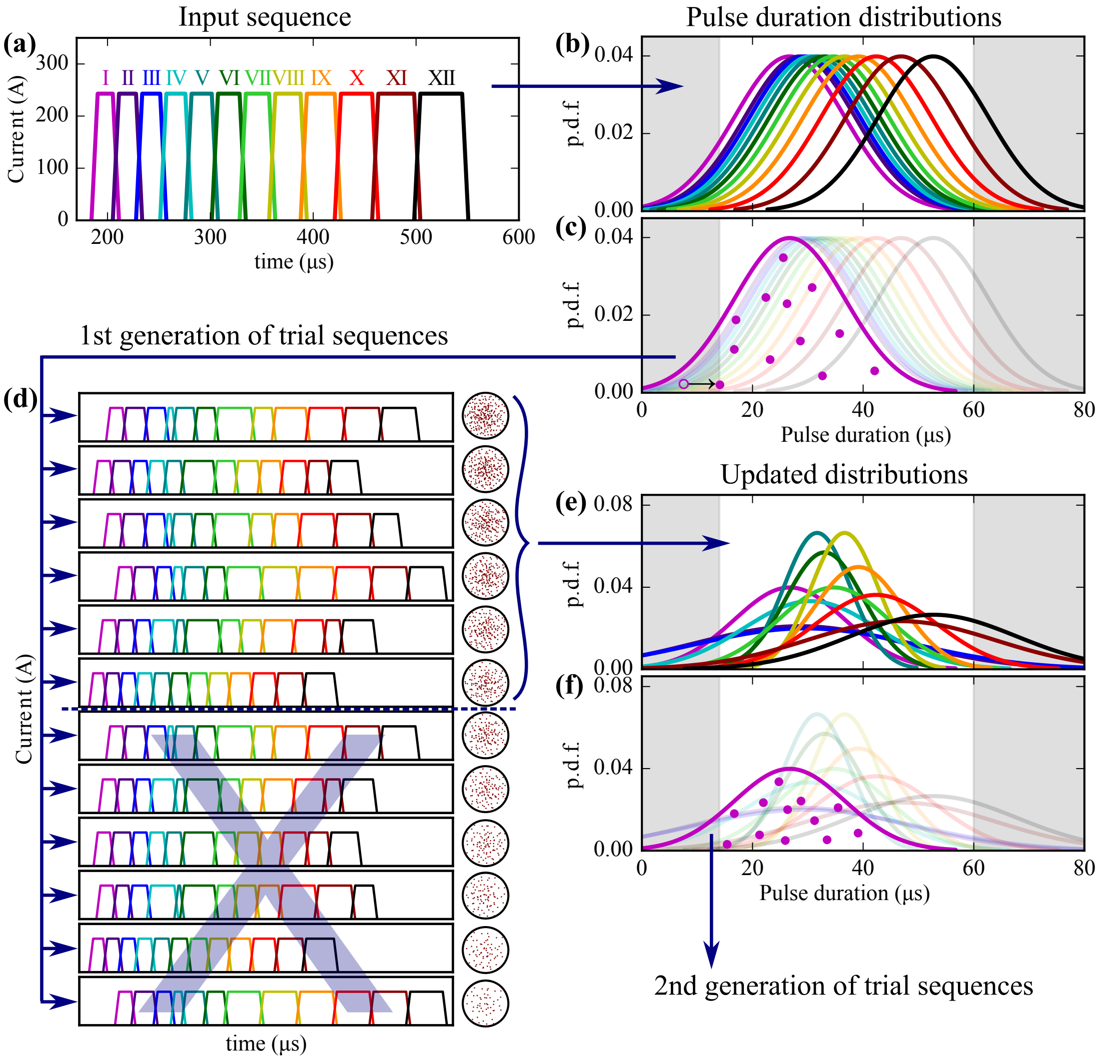

Figure 3.1: Schematic representation of the beginning of the CMA-ES optimisation: (a) an input deceleration sequence of current pulses as a function of time for each of the 12 coils is provided. (b) For each coil, a normal distribution is created, with mean corresponding to the pulse duration of the coil in the input sequence, and standard deviation of 10 μs. (c) A set of λ durations are randomly drawn from each distribution (shown for coil I only, purple dots; the vertical displacement of the purple dots is for illustrative purposes only). When a duration that is below (or above) the permitted pulse-length range is drawn (empty purple dot in the grey shaded region), it is replaced with the minimum (maximum) permitted value. (d) The first generation of λ candidate sequences is created from the randomly drawn durations. Each trial sequence is evaluated and ranked based on the number of particles with velocity within the target velocity range that it is simulated to yield (red dots within the black circles). The worst performing half of the sequences (below the blue dashed line) are discarded (blue shaded cross). (e) The distribution of durations for each coil is updated using the mean and standard deviation calculated from the best performing sequences. (f) λ new durations are drawn to create a second generation of trial sequences.

population of λ candidate sequences is obtained, where the pulse durations that compose each sequence are sampled from twelve multivariate normal distributions, one for each coil (Figure 3.1(b)-(d)). The mean and standard deviation of each of these distributions are computed from the preceding generation or, in the case of the first generation, from the durations of the standard pulse sequence that is being optimised with an arbitrarily selected standard deviation of 10 µs (Figure 3.1(b)). The constraints on pulse duration and switching times are implemented by replacing any pulse duration that exceeds (is lower than) the maximum (minimum) pulse duration with the maximum (minimum) permitted pulse length of 60 µs (14 µs) (Figure 3.1(c)); this method was found to be more efficient than simply discarding any sequences that fall outside the required constraints, or imposing a penalty proportional to the distance the candidate sequence lies outside the permissible region. The sequences thus generated are evaluated and ranked in fitness order based on the total number of particles decelerated to within the target velocity range (Figure 3.1(d)). The mean values and standard deviations of the multivariate distributions from which the durations are drawn are updated at every generation, and they are calculated from the weighted average of the best performing individual sequences (Figure 3.1(e)). Typical settings correspond to half of the sequences being discarded and the other half being used to update the distributions, with the weight of an individual sequence proportional to its ranking (Figure 3.1(d)). The updated mean values and standard deviations are then used with the covariance matrices to generate twelve new multivariate normal distributions from which λ sequences are created for the next generation, achieving the selection and recombination characteristics of an evolutionary algorithm (Figure 3.1(e),(f)).

The covariance matrix update is at the core of the CMA-ES method. An accurate estimate of the covariance matrix requires a large population of parameter sets. However, this comes at the expense of significantly increased computing time. A successful compromise, yielding faster convergence with high reliability, is achieved by using information from the best performing sequences in previous generations to update the covariance matrix. In this way, the best performing sequences of the first generation are used to create the second generation of sequences, with the worst performing sequences discarded. In essence (after the first generation), the distribution of pulse durations for each coil is established using the mean and standard deviation of the best performing sequences from the previous generation. This process of parameter optimisation means that the CMA method is nominally parameter-free, only requiring an initial choice of mean and standard deviation – values which are quickly adapted by the optimisation process. Throughout this iterative process, the globally best performing sequence is stored and replaced each time a better sequence is found. The optimisation process is deemed to have converged when the best performing sequence has not been updated for a given number of generations, chosen to be 50,000 generations in this work. The final best performing sequence is then selected and utilised experimentally.

Optimising for sequentially lower target velocities (below $v_z^{\text{final}} = 200\,\text{m}\,\text{s}^{-1}$), the nearest optimised evolutionary-algorithm sequence is taken as the starting sequence from which the new first generation is created. Thus, the evolutionary-algorithm-optimised sequence for $v_z^{\text{final}} = 200\,\text{m}\,\text{s}^{-1}$ is used as the input sequence for the target final velocity $v_z^{\text{final}} = 175\,\text{m}\,\text{s}^{-1}$, and in turn the optimised sequence for $v_z^{\text{final}} = 175\,\text{m}\,\text{s}^{-1}$ is used as the input sequence for the target final velocity $v_z^{\text{final}} = 150\,\text{m}\,\text{s}^{-1}$, and so on. The lowest target final velocity achievable is $v_z^{\text{final}} = 75\,\text{m}\,\text{s}^{-1}$, where the optimised sequence obtained (labelled "EA75" in Figure 3.2) uses the maximum permitted pulse duration for the last five coils. Attempts to realise lower final velocities fail to yield any particles in the target velocity range during the first generation of sequences and so cannot be optimised.

3.2.2 Comparison of sequences

As Figures 3.2 and 3.3 clearly illustrate, evolutionary-algorithm optimisation yields sequences that are markedly different from the ones generated by standard trajectory simulations considering a single synchronous particle. The standard sequence features monotonically-increasing current-pulse lengths applied to each coil; as the synchronous particle

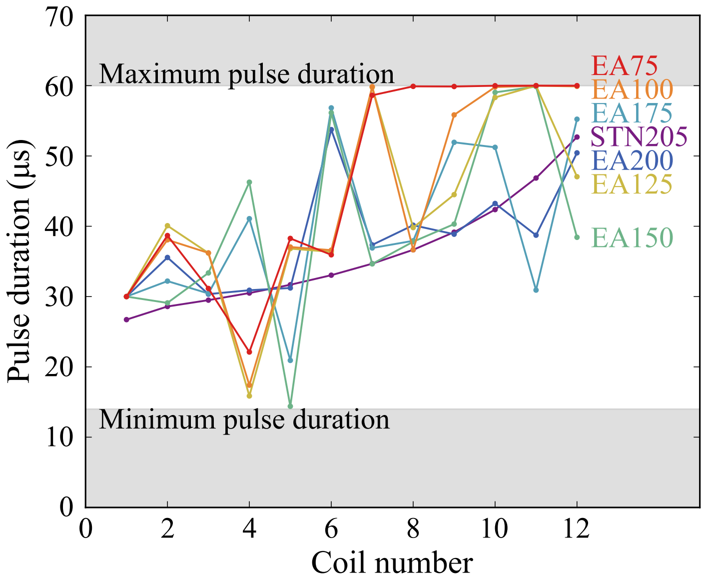

Figure 3.2: Pulse duration as a function of coil number for the standard sequence ($\phi_0 = 90°$, $v_z^{\text{final}} = 205\,\mathrm{m\,s^{-1}}$; labelled "STN205") and for six evolutionary-algorithm-optimised sequences (labelled "EA" followed by a number corresponding to the specific final velocity of each sequence). The lines are guides to the eye and the shaded regions correspond to pulse durations that are unfeasible with the current apparatus.

slows down, it takes longer to travel through each successive coil. The maximum experimentally feasible current-pulse duration and magnetic-field gradient across all coils imposes an upper limit on the final velocity that can be achieved with standard deceleration sequences; starting from an initial velocity of $500\,\mathrm{m\,s^{-1}}$, H atoms in the ground state can be decelerated to a minimum final velocity of $205\,\mathrm{m\,s^{-1}}$.

In contrast, sequences produced by evolutionary algorithms have pulse durations that vary in an oscillatory manner across consecutive coils. This results in some early coils being switched on for a longer time, enabling lower final velocities to be achieved compared to the standard sequence. Thus for the same constraints of minimum and maximum pulse length, and the same incoming beam, evolutionary-algorithm sequences can decelerate H atoms down to a velocity of $75\,\mathrm{m\,s^{-1}}$ (see Figure 3.3, bottom trace) – some $130\,\mathrm{m\,s^{-1}}$ lower than the minimum velocity achievable with the standard sequence. This is due to the fundamentally different approaches taken when generating the standard versus evolutionary-algorithm sequences. Essentially, the evolutionary strategy optimises sequences considering the properties of the entire beam and the passage of all particles

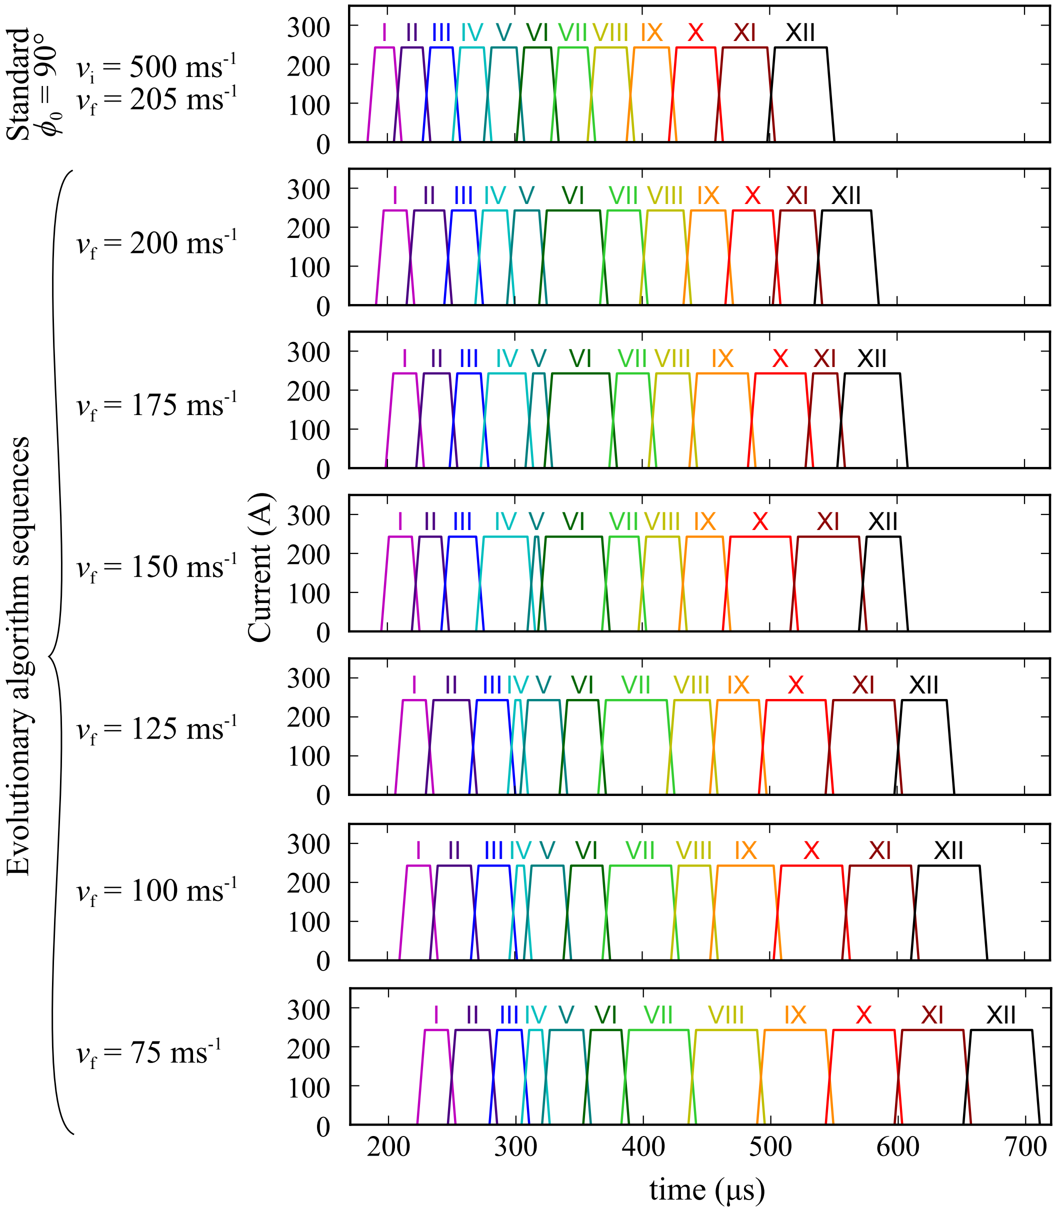

Figure 3.3: Schematic illustration of pulse sequences from standard trajectory simulations (first row) and from evolutionary algorithms (second to seventh row): the time profile of the current passed through each of the coils (labelled with roman numerals) is shown. The initial (v_i) and final (v_f) velocities of the sequences are annotated on the left.

through the decelerator, whereas the standard sequence considers the trajectory of only one synchronous particle.

3.3 Results

3.3.1 Experimental results

The current profiles generated by the standard sequence and a range of evolutionary-algorithm sequences (Figure 3.3) are implemented experimentally, with the resulting ToF traces presented in Figure 3.4. The experimental ToF profiles (solid lines) for the undecelerated (coils off) and decelerated (coils on) beams are displayed in red and blue, respectively. As expected, the fully decelerated peak (with the latest arrival time) in the decelerated traces progressively moves to longer ToF as the target velocity of the deceleration sequence is decreased. Importantly, the experimental ToF traces provide clear evidence that the optimised sequences are feasible; decelerated particles are detected at the anticipated ToF in all cases, even for velocities down to $75\,\mathrm{m\,s^{-1}}$. In all instances, the intensity and spread of the decelerated peak is reproduced by the simulations (shaded profiles).

For each deceleration sequence, ToF traces are recorded at five different detection distances (as indicated by the series of REMPI detection laser positions in Figure 2.3(b)) to facilitate an independent verification of the velocity of H atoms contributing to the final peak. One such series of ToF profiles at a range of detection distances is provided in Figure 3.5, corresponding to the evolutionary-algorithm sequence with a target final velocity of $200\,\mathrm{m\,s^{-1}}$. By plotting the change in detection distance against the change in ToF (schematically represented by the dashed green line across the panels in Figure 3.5), the mean experimental velocity of the particles contributing to the decelerated peak can be estimated. With this approach, the velocity is experimentally determined to be $210 \pm 10\,\mathrm{m\,s^{-1}}$, in agreement (within error) with the target velocity of $200 \pm 10\,\mathrm{m\,s^{-1}}$ utilised in the optimisation procedure. The slight overestimate in the experimental velocity established with this simple approach is attributed to the presence of faster-moving particles in the decelerated peak, as shown in Figure 3.7 and discussed below.

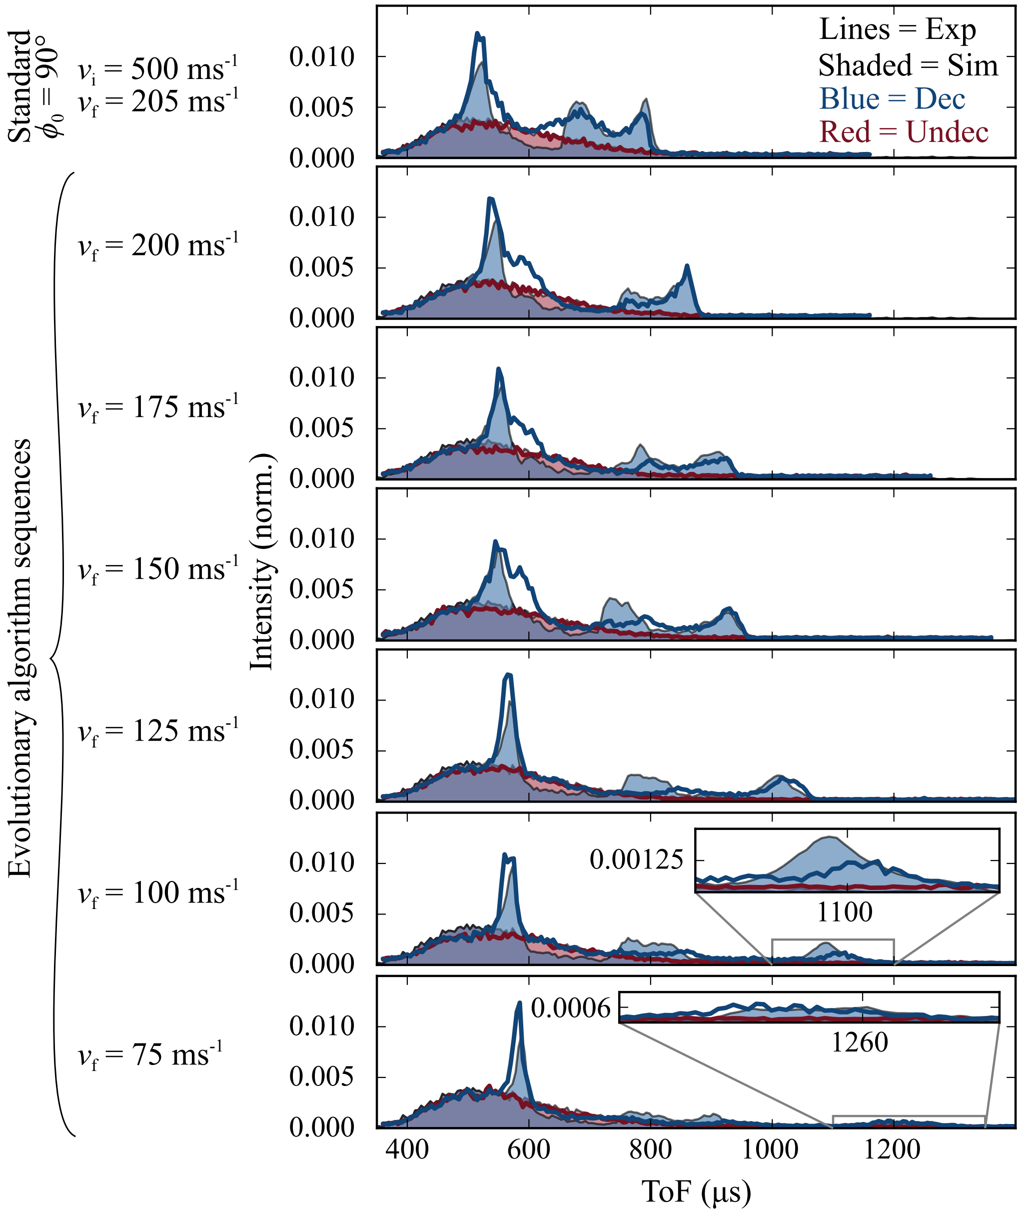

Figure 3.4: Time-of-flight (ToF) traces from standard (first row) and evolutionary-algorithms-optimised (second to seventh row) sequences (detection distance d = 260 mm). The initial (v_i) and final (v_f) velocities of the sequences are annotated on the left. Solid lines indicate the experimental traces and shaded areas show the simulated traces. For each ToF profile, both the undecelerated (magnetic fields off, red colour) and decelerated (magnetic fields on, blue colour) profiles are shown.

The simulation analysis enables one to disentangle the contribution that decelerated H atoms make to the final ToF distribution (Figure 3.6). This confirms that state selectivity of the particles in the decelerated bunch is attained, and that the final desired velocity is achieved. The analysis is illustrated for the case of the evolutionary-algorithm-generated

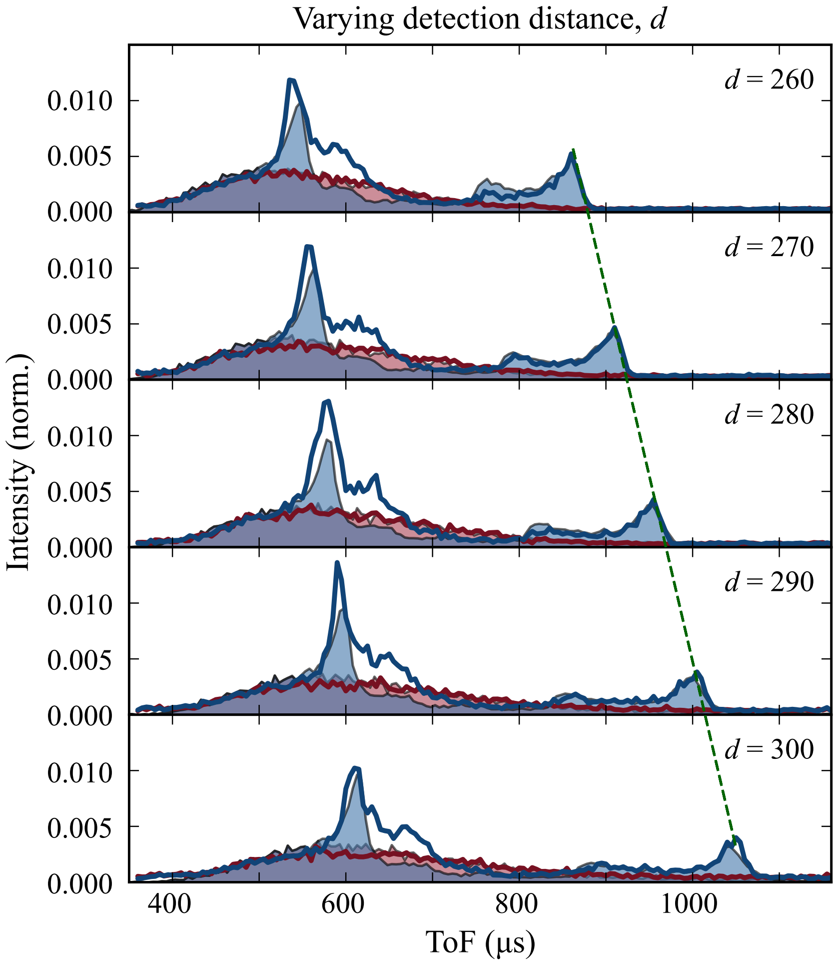

Figure 3.5: ToF traces obtained at 5 different detection distances (see Figure 2.3) are shown for the evolutionary-algorithm sequence yielding a final velocity of 200 m s^{-1}. The green dashed line illustrates the delay in the decelerated signal as a function of detection distance, from which the velocity of the decelerated peak can be calculated.

sequence for a final velocity of 200 ± 10 m s^{-1}. The panels on the left hand side column of Figure 3.6 plot the final velocity of the simulated particles against their ToF – shown for the undecelerated (first row) and decelerated (second row) beams, with the contribution from LFS (third row) and HFS (fourth row) particles explicitly displayed. From the "Dec (LFS)" plot, it can be seen that the majority of the particles in the decelerated peak have a ToF arrival time of approximately 830 µs, and that this packet of decelerated LFS particles travels with a final velocity centred on 200 m s^{-1}. A minor contribution to the decelerated peak arises from the small number of particles (in both LFS and HFS states) with an initial velocity around 300 m s^{-1} that are not decelerated by the magnetic fields.

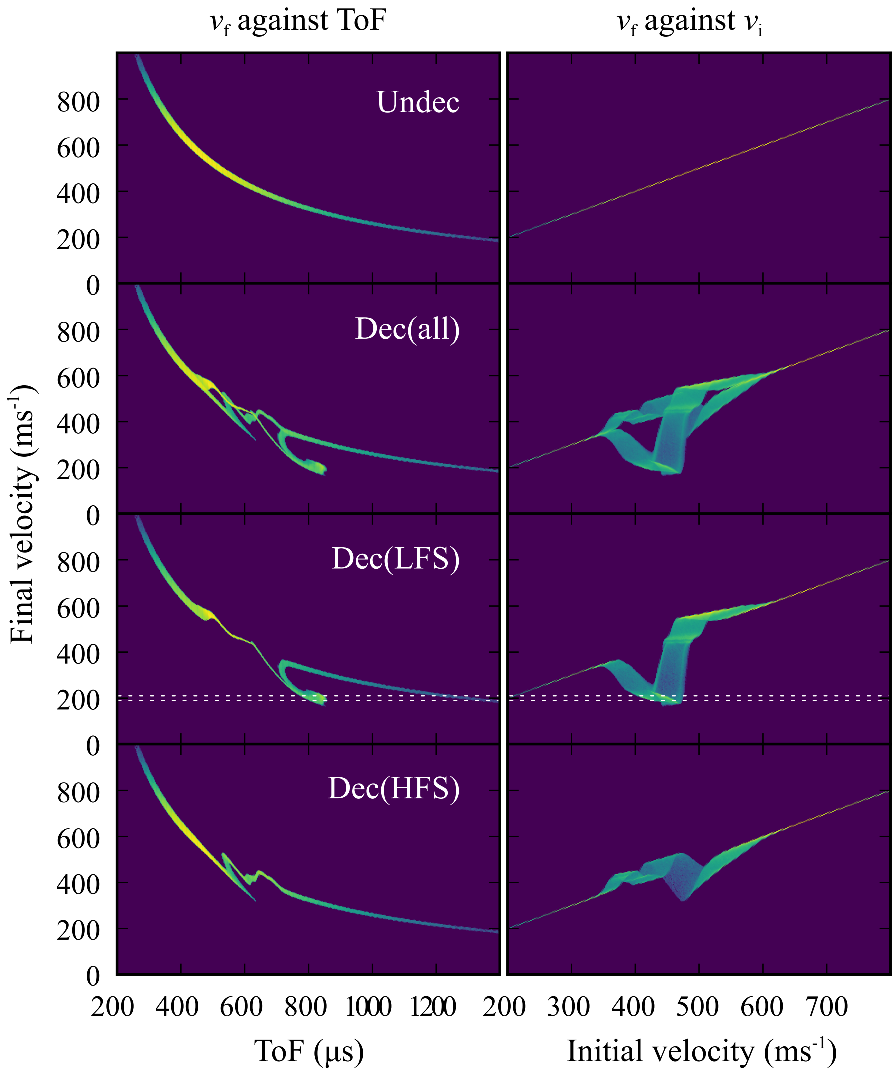

Figure 3.6: Simulation analysis for the $v_f = 200\,\mathrm{m\,s^{-1}}$ sequence, depicting the final velocity of the particles plotted as a function of ToF (left panels) and initial velocity (right panels) for the various components of the beam. This includes the undecelerated beam "Undec", the decelerated beam including all quantum states "Dec (all)", the decelerated beam including only LFS states "Dec (LFS)", and the decelerated beam including only HFS states "Dec (HFS)". The horizontal white dotted lines in the two "Dec (LFS)" panels indicate the target velocity range.

Thus it should be noted that – while all particles travelling at approximately $200\,\mathrm{m\,s^{-1}}$ are in LFS quantum states – not all particles forming part of the "decelerated" peak in the ToF trace have been decelerated, as there is a minor contribution from initially slow-moving particles in the incoming beam (see also Figure 3.7). The panels on the right of Figure 3.6 show plots of the final velocity of the simulated particles against their initial velocity.

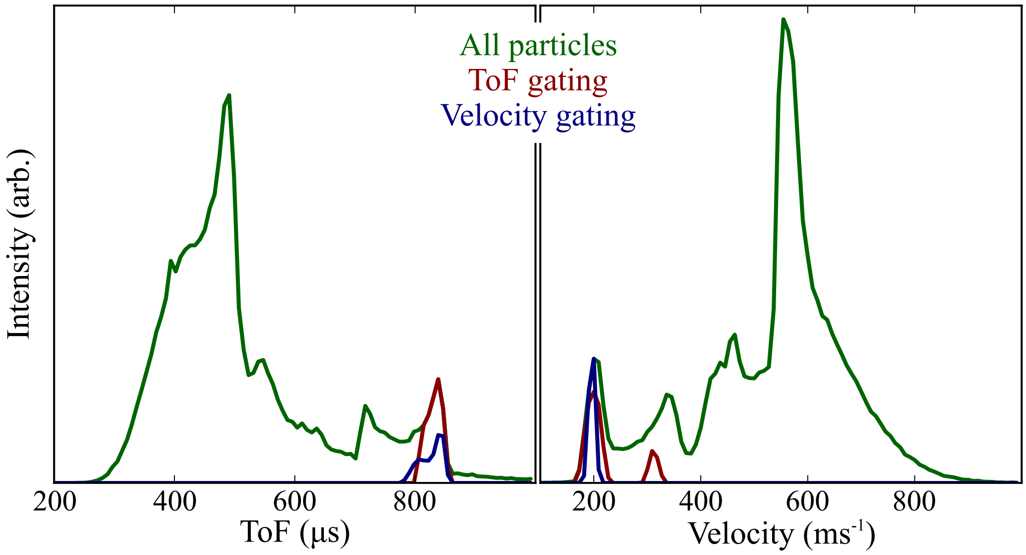

Figure 3.7: Simulated ToF trace (left panel) and the corresponding velocity distribution (right panel) for a beam of hydrogen atoms after deceleration with the evolutionary-algorithm sequence yielding a final velocity of 200 m s^{-1} (green traces). The particles with velocity between 190 and 210 ms^{-1} arrive at similar ToF, resulting in a single peak around 830 µs (blue traces). However, the particles arriving at ToF between 810 and 860 µs (the decelerated ToF peak) travel at two distinct velocities, centred around 200 and 300 ms^{-1} (red traces). Thus, whilst all particles travelling at the target velocity belong to the decelerated peak in the ToF, not all particles belonging to that peak are decelerated.

Here, it can be seen that all particles that are decelerated to a final velocity of approximately 200 m s^{-1} are LFS particles that started with an initial velocity between 400 and 480 m s^{-1}. Repeating the same analysis for all sequences confirms that the desired final velocity is reached in all cases.

3.3.2 Simulation analysis

Evolutionary algorithms can achieve greater deceleration than would be possible using standard switching sequences, given the experimental constraints. In order to understand how this is achieved, simulations of the standard pulse sequence yielding a final velocity $v_z^{\text{final}} = 205\,\text{m s}^{-1}$ are compared with those of the optimised evolutionary-algorithm sequence with a final velocity $v_z^{\text{final}} = 200\,\text{m s}^{-1}$ (the top two sequences in Figure 3.3, respectively). Figure 3.8 shows the phase-space plots of the particles contributing to the

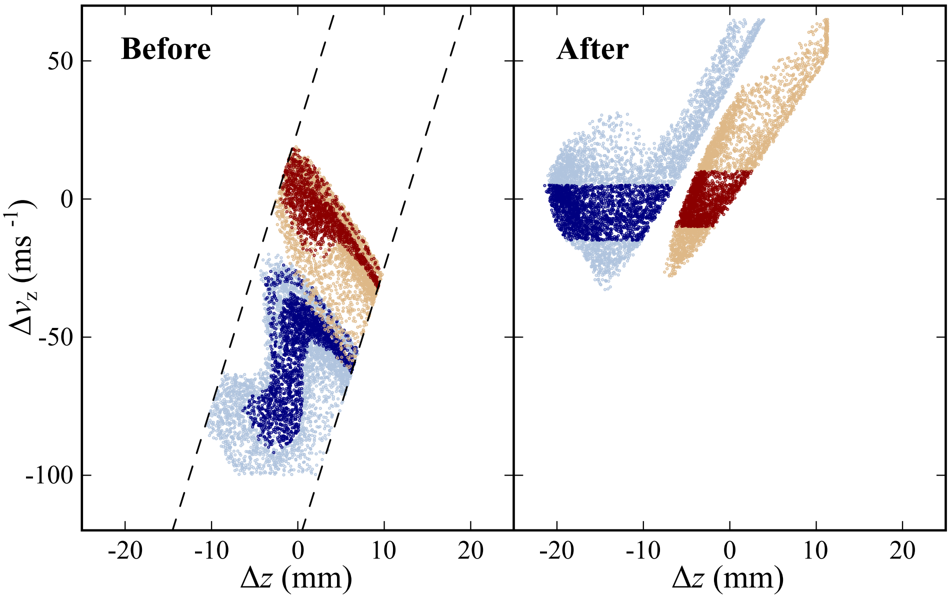

Figure 3.8: Phase-space plots of the difference in forward velocity against the difference in position for the simulated particles with respect to the synchronous particle located at $(\Delta x, \Delta y, \Delta z = 0,0,0)$ for the standard sequence ($\phi_0 = 90°$, $v_z^{\text{final}} = 205\,\text{m}\,\text{s}^{-1}$; red) and the evolutionary-algorithm sequence ($v_z^{\text{final}} = 200\,\text{m}\,\text{s}^{-1}$; blue) before (left panel) and after (right panel) deceleration. Only the particles in the initial beam that are decelerated to velocities within the relevant ranges are shown. Particles decelerated to within $\pm 10\,\text{m}\,\text{s}^{-1}$ of the target velocity are indicated by dark blue and dark red dots for the evolutionary-algorithm and standard sequences, respectively. A wider contour of particles decelerated to the target velocity $\pm 70\,\text{m}\,\text{s}^{-1}$ is also shown for reference, indicated by light blue and light red dots for the evolutionary-algorithm and standard sequences, respectively. The dashed lines in the left panel represent the boundaries of the initial phase-space distribution of the undecelerated beam. Note that the dashed lines are not vertical due to the initial expansion of the incoming beam before it reaches the first decelerator coil (with faster particles travelling further than slower particles).

decelerated peak, both before entering the decelerator (left panel) and after all particles in the phase-stable region have exited the final coil (right panel). The standard sequence behaves as expected: as it is based on the trajectory of a single synchronous particle with an initial velocity of $500\,\text{m}\,\text{s}^{-1}$, particles with a similar velocity and position to the synchronous particle also fall within the phase-stable region. In contrast, the optimised evolutionary-algorithm sequence considers an ensemble of particles with a larger distribution of velocities and positions. It preferentially addresses the slower-moving particles in the initial distribution, meaning that less translational energy needs to be removed to achieve the same final velocity. The particles decelerated by the evolutionary-algorithm sequence emerge from the decelerator with a similar spread in velocity but a larger spread in position than those decelerated by the standard sequence (Figure 3.8, right panel). This is further illustrated in Figure 3.9, where it can be seen that the velocity (a) and spatial (c) distributions in the x (and also y, by symmetry) dimension are comparable between the two sequences (both before and after deceleration) but there are differences in the z dimension. Whilst the spread in z velocity decreases during the deceleration (b), the spread in z position increases (d). The evolutionary-algorithm sequence yields decelerated particles with a 14 mm spread of positions along the z axis, whereas the standard sequence distribution along z is only 8 mm.

Ultimately, the aim is to combine the decelerator with an ion trap to undertake collision studies between radical species and trapped ions (see section 1.3). Provided the distribution in the xy plane matches the acceptance of the ion trap – which can be achieved through the inclusion of additional focusing components (see chapter 4) – the spread in z is less important to the intended application than the number of particles that reach the reaction region. To this end, the evolutionary-algorithm sequence far exceeds the performance of the standard sequence: the number of decelerated particles is 40% higher for the optimised sequence for a comparable final velocity.

The success of the evolutionary-algorithm sequence can be attributed to its increased phase-space acceptance and to its ability to address particles that are already travelling slower and therefore require less deceleration. The latter is observed across all of the evolutionary-algorithm sequences studied; the lower the target velocity to be optimised for, the lower the initial velocity of the particles that are decelerated. Qualitatively, this can be seen in the progressively delayed switch-on time of the current pulse applied to the first coil in Figure 3.3. To achieve lower final velocities using a standard sequence, a synchronous particle with a lower initial velocity could be selected, provided a lower effective phase angle is also employed to ensure that the maximum pulse duration does not exceed 60 μs (for the decelerator used in this work). A synchronous particle with an initial velocity of, for example, $450\,\mathrm{m\,s^{-1}}$ – comparable to the modal velocity selected by the evolutionary-algorithm-optimised sequence yielding the same final velocity – could

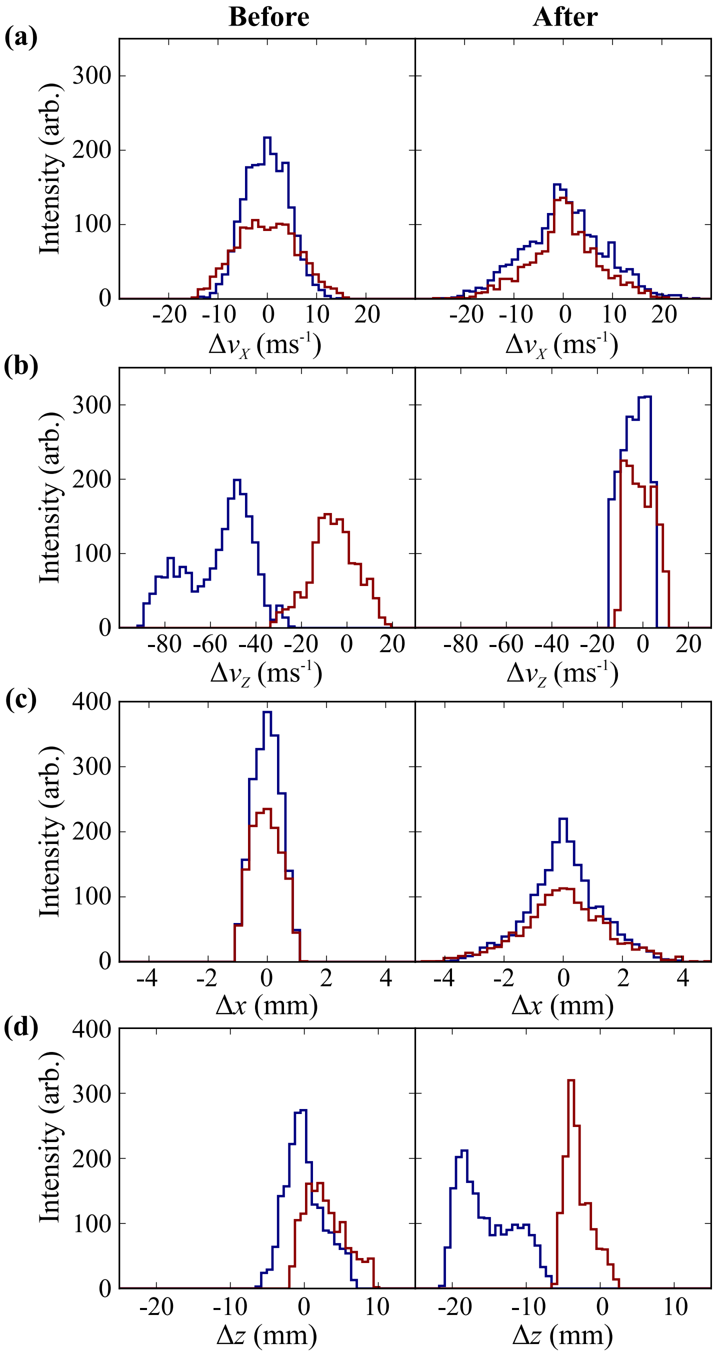

Figure 3.9: Histograms of the x velocity (a), z velocity (b), x position (c) and z position (d) of the simulated particles that are decelerated to within $\pm 10\,\mathrm{m\,s^{-1}}$ of the target velocity. Particles are plotted with respect to the synchronous particle located at $(\Delta x, \Delta y, \Delta z = 0,0,0)$ with velocity $(\Delta v_x, \Delta v_y, \Delta v_z = 0,0,0)$, before (left column) and after (right column) deceleration using the standard sequence ($\phi_0 = 90°$, $v_z^{\mathrm{final}} = 205\,\mathrm{m\,s^{-1}}$; red) or the evolutionary-algorithm sequence ($v_z^{\mathrm{final}} = 200\,\mathrm{m\,s^{-1}}$; blue). Note that, due to the cylindrical symmetry of the decelerator, $\Delta v_x = \Delta v_y$ and $\Delta x = \Delta y$.

be decelerated to $200\,\mathrm{m\,s^{-1}}$ using a standard sequence, with $\phi_0 = 53°$. However, simulations demonstrate that this approach yields significantly fewer decelerated particles. While one might anticipate an increase in the number of decelerated particles transmitted through the decelerator when the phase angle is reduced (see section 2.1), this effect is outweighed by the decrease in the number of particles in the initial beam that have a low enough velocity to be addressed by the sequence. Indeed, for the same final velocity, the standard sequence with $\phi_0 = 53°$ ($v_z^{\mathrm{initial}} = 450\,\mathrm{m\,s^{-1}}$) yields 50% fewer decelerated particles compared to the standard sequence with $\phi_0 = 90°$ ($v_z^{\mathrm{initial}} = 500\,\mathrm{m\,s^{-1}}$) and 65% fewer particles than the evolutionary-algorithm-optimised sequence. Thus, employing standard sequences with a lower phase angle and lower initial velocity synchronous particle results in significantly fewer decelerated particles. Such sequences cannot be designed to reach lower final velocities than standard sequences with $\phi_0 = 90°$, unless longer pulse durations are used.

The velocity profile of decelerated particles as a function of time also looks distinctly different for the standard and evolutionary-algorithm sequences (see Figure 3.10). For the standard sequence, the packet of decelerated atoms is kept together in phase space and experiences a periodic magnetic-field pattern and thus a stepwise decrease in mean velocity at each stage. For the optimised sequence, on the other hand, the deceleration appears to occur in two stages: in the first stage (coils 1 to 5), particles that are initially moving faster are targeted, whilst the velocities of the slower particles are barely affected by the magnetic fields at all. This results in the spread of velocities that can be observed just before coil 6 is switched off (sixth vertical yellow line). In the second stage, all particles in the phase-stable bunch are decelerated by the remaining coils and v_z gets re-focused, as shown by the increased density of particles around $v_z = 200\,\mathrm{ms^{-1}}$ at the end of the deceleration process. In order to keep the entire packet together and refocus the velocity spread, the initially slower particles (which are now the faster particles in the packet, owing to the deceleration of the previously faster particles) are preferentially decelerated at a higher rate. The velocities of all particles trend downwards for each successive coil from coil 6 to the end of the decelerator. While this description of the two-

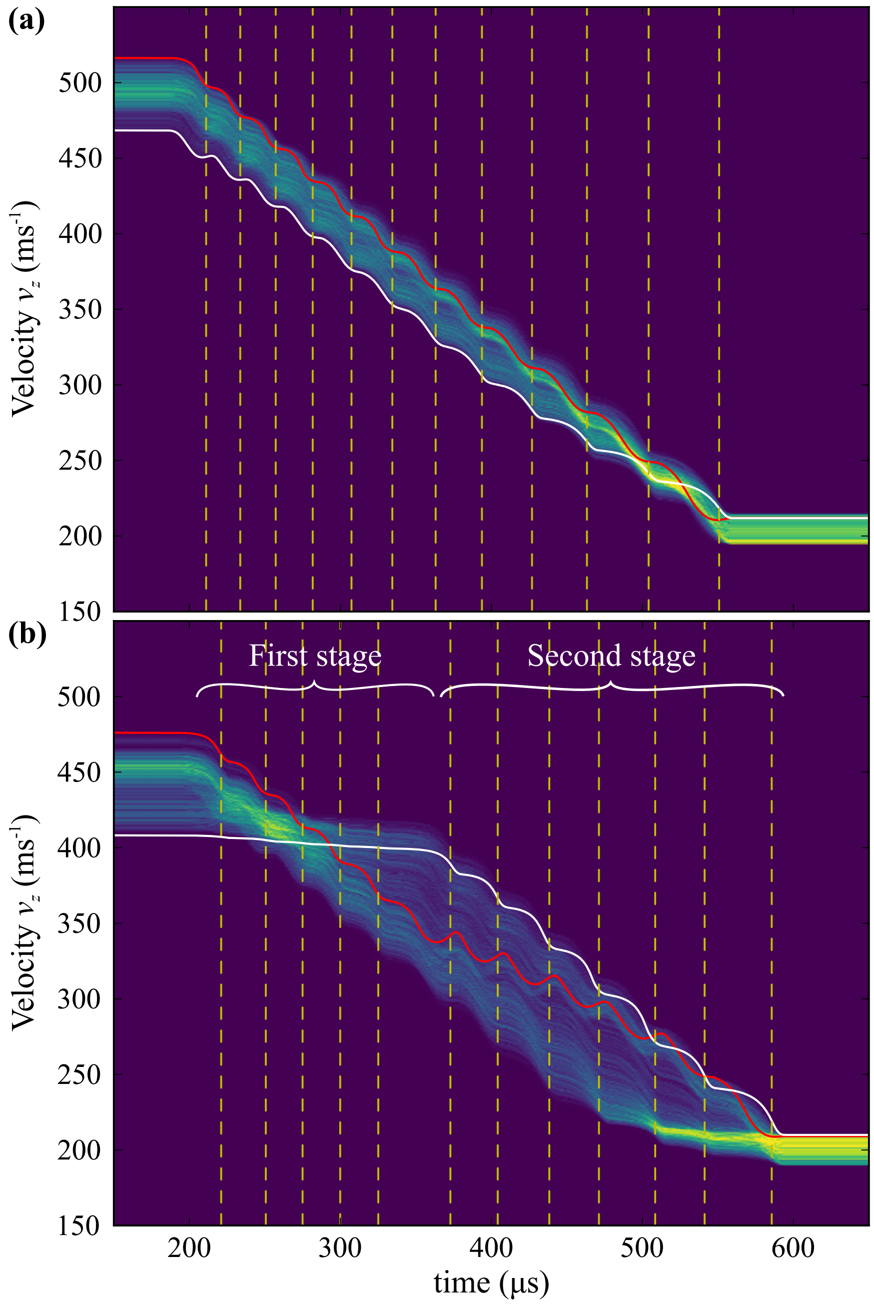

Figure 3.10: Velocity distribution of the decelerated packet as a function of time for the standard sequence ($\phi_0 = 90°$, $v_z^{\text{final}} = 205\,\text{m}\,\text{s}^{-1}$; (a)) and the evolutionary-algorithm sequence ($v_z^{\text{final}} = 200\,\text{m}\,\text{s}^{-1}$; (b)), with intensity ranging from low (purple) to high (yellow). The vertical yellow lines denote the switch-off time of each coil. The red and white lines trace the trajectories of a pair of particles with fast and slow initial velocities, respectively. Note that only particles that are successfully decelerated to within $\pm 10\,\text{m}\,\text{s}^{-1}$ of the target velocity are shown.

stage deceleration is an oversimplification, explicitly considering only the extremes of the initial velocity distribution, it helps to illustrate the differences in how the evolutionary-

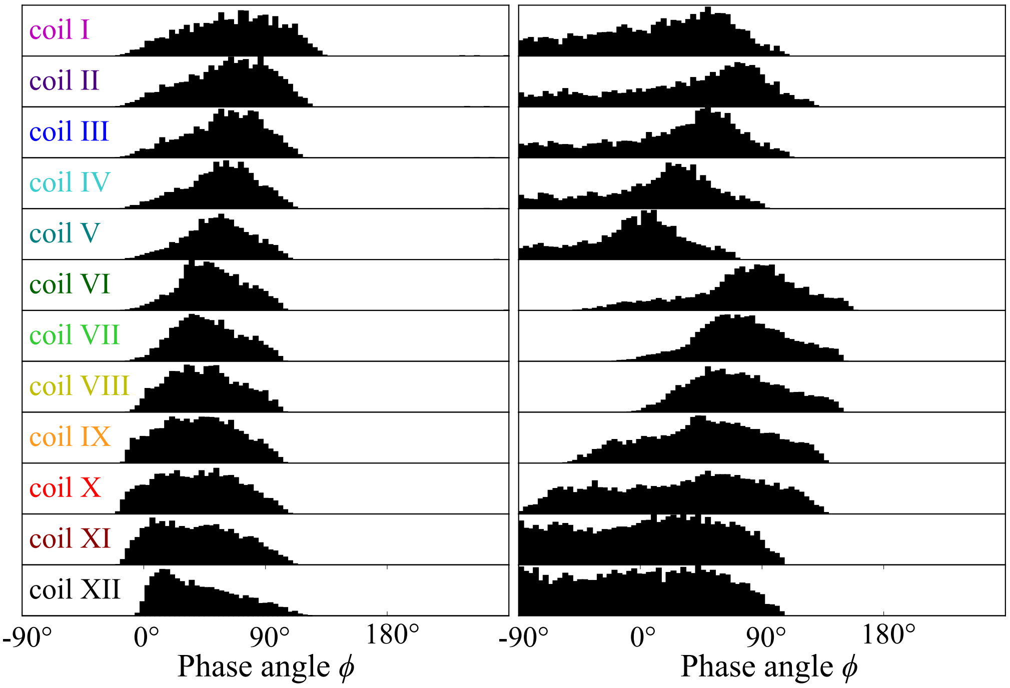

Figure 3.11: Histograms of the phase angle of all particles that are successfully decelerated to within ±30 $\mathrm{m\,s^{-1}}$ of the target velocity, at the instant of switch-off for each of the coils (I to XII) for the standard sequence ($\phi_0 = 90°$, $v_z^{\text{final}} = 205\,\mathrm{m\,s^{-1}}$; left) and the evolutionary-algorithm sequence ($v_z^{\text{final}} = 200\,\mathrm{m\,s^{-1}}$; right).

algorithm and standard sequences achieve deceleration of the phase-stable packet, and the way that the evolutionary algorithm is able to capture a broader range of the initial phase-space distribution.

The two-stage deceleration process adopted by the evolutionary-algorithm sequence can also be observed in the distribution of phase angles exhibited by all particles in the decelerated packet at the switch-off time of each coil (see Figure 3.11), corresponding to the positions of particles in the coil at the time when the magnetic field is switched off. The range of phase angles in particles addressed by the standard sequence (left column) is fairly uniform and narrow for all coils, with the peak in the distribution a little below $\phi = 90°$ for the earlier coils and slowly shifting to lower phase angles for progressively later coils. In the evolutionary-algorithm sequence (right column), there is a clear discontinuity between coils 5 and 6, the point at which the first stage of deceleration (outlined above) ends and the second stage begins. In the first stage, corresponding to coils 1-5, the distributions peak before $\phi = 90°$ and this peak again shifts progressively to lower phase angles (as seen in the standard sequence, but at a faster rate). In contrast to the standard sequence, the distributions are much broader due to the slower particles in the bunch that are lagging behind the faster ones. In the second stage, from coil 6, an appreciable number of particles have travelled further than the middle of the coil ($\phi = 90°$) by the time the coil is switched off and, as a result, they get decelerated and then slightly accelerated again. As close inspection of Figure 3.10 reveals, some of the initially fastest-moving particles are periodically decelerated and then accelerated as they pass through coils 6-12, although there is still a net deceleration of all particles as they pass through each of the coils in the second stage. This unintuitive approach serves to bunch the velocity distribution, returning the desired narrow distribution of velocities at the exit of the decelerator.

3.3.3 Performance of the evolutionary-algorithm sequences

Optimised evolutionary-algorithm sequences outperform standard sequences due to two key factors. Firstly, the optimised sequences can circumvent the issues faced by the standard sequences by targeting particles that are already travelling slower without needing pulses that are too long to be implemented experimentally. Given that $500\,\mathrm{m\,s^{-1}}$ is the most probable velocity of particles in the beam entering the decelerator, addressing particles with $v_z^{\mathrm{initial}} < 500\,\mathrm{m\,s^{-1}}$ means that fewer particles are available for deceleration. This trade-off between addressing already slower particles and maintaining the highest possible number of decelerated particles is automatically optimised by the algorithm. Secondly, the evolutionary sequences can target a larger spread of initial velocities since the optimisation is performed by considering a large number of particles with a range of initial positions and velocities instead of one synchronous particle, with focusing effects automatically accounted for. As a result, the evolutionary-algorithm sequences have a larger phase-space acceptance and achieve significantly enhanced numbers of decelerated particles than standard sequences, for a comparable final velocity.

The beauty of using evolutionary algorithms is that one can tweak the criteria against which the pulse-sequence parameters are optimised. Should a narrower distribution in phase space be required, one could choose to optimise the sequences both in terms of final velocity and spread of position by adding an additional criterion to the fitness function. However, the phase-space density cannot be increased during the deceleration, since the forces acting on the particles are conservative. Therefore, such optimisation would come at the expense of a reduced number of particles. Alternatively, an optimisation that maximises the number of particles with a specific longitudinal velocity arriving at a specific longitudinal position might result in an increased transverse velocity and transverse position spread of the decelerated particles. This is a key difference between the optimised evolutionary-algorithm sequences and the standard sequences. Typically, deceleration sequences can be designed to rotate the phase-space distribution to minimise the spread in z, or to compact the longitudinal velocity distribution at the expense of the distribution in z; it is difficult to devise standard sequences that minimise the spread in both longitudinal velocity and space. Through the use of evolutionary algorithms, one could select a larger part of the phase-space distribution of the incoming beam in order to optimise both parameters simultaneously. The two-stage deceleration approach seen with the evolutionary-algorithm sequence allows species to enter the phase-stable region as they move through the apparatus, increasing the number of particles addressed throughout the second stage of deceleration such that the total number of species decelerated to within the target velocity range is maximised at the end of the decelerator.

The inclusion of additional criteria was not explored in this implementation, for the dual reasons of wanting to keep the optimisation as simple as possible and because a narrow spatial distribution (at least along the z axis) is not necessary for the intended applications described in this thesis. Additional focusing elements positioned after the decelerator, as discussed in chapter 4, permit focusing of the decelerated packet of particles for reaction studies or other applications. Indeed, there are no limits (in principle) to the criteria against which the parameters can be optimised using evolutionary algorithms; provided that at least one particle meets the criteria during the first generation of sequences, the evolutionary-strategy approach returns the best sequence for a given set of restraints.

Computational time could potentially become a limiting factor for the optimisation of very complicated systems. The optimisation procedure described here takes on the order of 1 hour per sequence, compared to only seconds required to generate a standard sequence, performed on a quad-core processor.

3.4 Conclusion

Optimised sequences of current pulses applied to a 12-coil Zeeman decelerator have been generated using a CMA-ES approach, for a range of target final velocities. Experimental measurements confirm that sequences optimised using an evolutionary algorithm yield both higher numbers of decelerated particles (for the same final velocity) and enable lower final velocities to be reached than is possible with standard sequences. Specifically, optimised sequences achieve a 40% increase in the number of decelerated particles for a final velocity of $200\,\mathrm{m\,s^{-1}}$, and they can produce particles with a modal kinetic energy that is 98% lower than that of the incoming beam, removing significantly more kinetic energy than can be achieved with standard sequences. These results are obtained using the same incoming beam and by imposing the same set of experimental restrictions for coil current (243 A), repetition rate (10 Hz) and maximum pulse duration (60 µs). The experimental ToF traces are well reproduced by three-dimensional particle-trajectory simulations, which reveal some surprising differences in how deceleration is achieved in the standard and optimised approaches.

Standard sequences consider the passage of a single synchronous particle travelling with (typically) the modal initial velocity of the incoming beam. In contrast, optimised sequences target particles that are travelling slower than the synchronous particle – and therefore require less deceleration to reach the same final velocity – while at the same time compensating for the loss of intensity in the beam by addressing a wider range of initial velocities. By selecting progressively slower portions of the initial distribution, optimised sequences can decelerate particles to lower final velocities without exceeding the maximum pulse duration. This successful deceleration of particles with a range of initial velocities is achieved by a non-trivial deceleration mechanism that could not have been conceived using a standard optimisation method.

In conclusion, using evolutionary algorithms to optimise deceleration sequences leads to significant increases in both the number of decelerated particles and the amount of deceleration achievable with a 12-stage Zeeman decelerator. The approach described here is highly versatile, robust and easily implemented. Employing a CMA-ES optimisation method could greatly improve the efficiency and operation of similar devices, potentially circumventing the need for longer decelerators in order to access lower final velocities. For instance, the deceleration of metastable N atoms reported in section 2.4 could be significantly improved by optimising the deceleration sequences using evolutionary algorithms. Besides obtaining a larger number of decelerated particles to make the decelerated "shoulder" more visible above the undecelerated distribution, it might be possible to remove more kinetic energy without modifying the apparatus. This could be achieved by varying the pulse durations – as demonstrated in this chapter – or even optimising the pulse amplitudes for the individual coils, although this would be more challenging to implement experimentally. Indeed, the CMA-ES method could be used to optimise sequences for a wide range of desired outputs by modifying the optimisation criteria (i.e. fitness function) of the algorithm. For example, the number of decelerated particles arriving at a specific position within a given velocity range could be maximised whilst the number of particles with velocities outside the target range at the desired position is simultaneously minimised, achieving both deceleration and spatial purification of the beam at a specific post-decelerator position. Optimisation using evolutionary algorithms is extremely powerful in generating superior deceleration sequences with improved abilities to decelerate more particles and reach lower velocities. The use of such evolutionary strategies could lead to significant advances in controlling the properties of cold beams formed by the manipulation of magnetic or electric fields.

4 | Magnetic guiding

4.1 Principles of magnetic guiding

The beam emerging from a Zeeman decelerator, as described in chapter 3, is typically composed of a mixture of decelerated, partially decelerated and undecelerated particles, in addition to precursor molecules and seed gas particles. In order to use a Zeeman-decelerated beam to study radical reactions with controlled collision energy and chemical composition (see section 1.3), the beam must undergo further manipulation prior to entering the interaction region. In this chapter, a magnetic guide is employed to purify the post-deceleration beam, such that only the decelerated radicals reach the interaction region for collisional studies. The requirements of such a guide are that it only transmits a beam of pure radicals with a narrow velocity distribution, and that this velocity is tuneable to facilitate the study of, for example, ion–radical reactions as a function of collision energy. The combination of a Zeeman decelerator with a linear Paul ion trap offers a highly promising approach for the study of a diverse range of ion–radical reactions under cold, controlled conditions, provided that only the decelerated radical species of interest are transmitted into the ion trap region.

Pure radical beams have been produced by photo-detachment of an electron from negative ions following mass and velocity selection [103]. However, the applicability of this method is limited to species that readily accept a negative charge, and does not allow for the purification of an existing radical beam. A different approach consists in coupling magnetic or electric deflection with mechanical barriers [104]. Straight magnetic and electrostatic deflectors have successfully been employed to spatially separate the *ortho* from the *para* spin isomer of water [105] and the *cis* from the *trans* conformer of 3-aminophenol [106], respectively. A pure beam of state-selected NO radicals has been obtained by combining an electrostatic hexapole state-selector with a beamstop and a diaphragm [107]. In such an approach, a 2 mm beamstop is placed in the middle of the

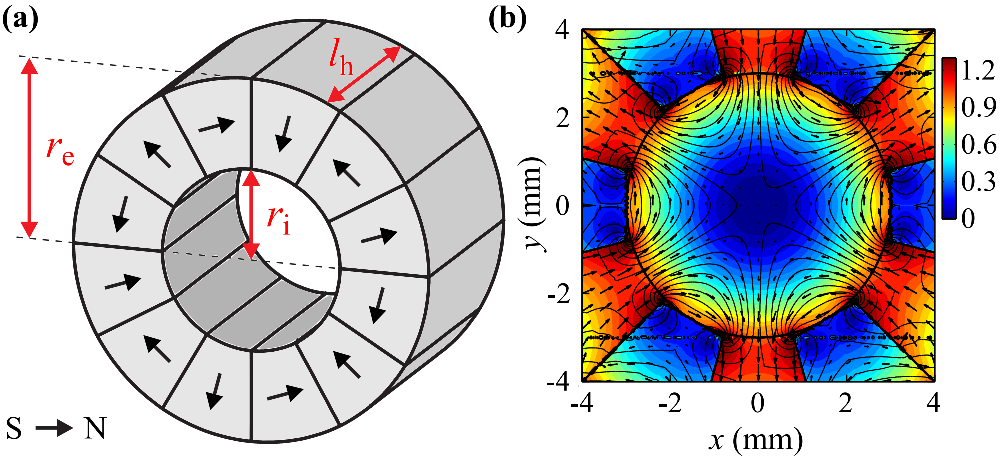

Figure 4.1: (a) Schematic representation of one Halbach array of permanent magnets in hexapolar configuration. (b) Magnetic-field magnitude (in Tesla) in the xy plane for an Halbach array with remanence $B_0 = 1.3\,\mathrm{T}$, internal radius $r_i = 3\,\mathrm{mm}$, external radius $r_e = 6\,\mathrm{mm}$ and thickness $l_h = 7\,\mathrm{mm}$.

hexapole – where the radicals of interest are most de-focused – and is centred on the beam axis, to block the carrier gas; a 2 mm-diameter diaphragm is placed after the hexapole to remove any remaining carrier gas. Although very high state purity (99%) is achieved, fine manipulation of the velocity distribution of the NO radicals is not feasible. Bent magnetic guides consisting of a series of permanent hexapole or octopole magnets have been shown to guide neutral beams of metastable neon [108] and lithium [53] atoms respectively, and merged beam experiments have very successfully employed magnetic benders to direct paramagnetic particles onto the same axis as a beam of polar molecules [55, 56] (see section 1.2.2). A design for a bent magnetic guide composed of hexapole magnets combined with current carrying wires has been proposed [109], where permanent hexapole magnets in a Halbach array configuration [110, 111] (Figure 4.1(a)) serve to deflect the decelerated particles away from the beam axis whilst current carrying wires enhance the velocity resolution and quantum-state selectivity of the guided particles. The zero magnetic field at the centre of the arrays (Figure 4.1(b)) is optimal for guiding particles in low-field-seeking quantum states, as such particles minimise their Zeeman energy by travelling through the centre of the arrays. Hexapole magnets in a Halbach array configuration are particularly suitable here due to their characteristic strong magnetic fields within the cylinder and weak magnetic fields outside it, which minimises interference with other electromagnetic components of the experiment. Similar arrays have been employed to improve the transverse confinement of particles in between Zeeman-deceleration stages [76], and are proposed in a design for a hydrogen storage ring [59].

Taking inspiration from previous work [109], a novel magnetic guide is proposed here: using Halbach arrays to deflect and collimate the decelerated particles off the main beam axis and around a blade that physically stops any faster or undeflected species travelling on-axis, before returning the decelerated species back to the main beam axis. Such a design increases both the practicality and flexibility of the guide. Indeed, returning the decelerated particles to the beam axis is necessary to maintain a collinear arrangement between the decelerator and the ion trap. This, besides simplifying alignment, circumvents the need to move the ion trap target when different velocities are guided. While originally developed to serve as a filter for Zeeman-decelerated radical beams, the magnetic guide reported in this chapter is immensely versatile. It could, for example, operate as a stand-alone radical beam filter – producing radical beams with a tuneable and narrow velocity distribution from a collimated effusive source – making the device relevant for a wide range of potential applications.

4.2 Deflection studies with one Halbach array

Initially, the beam deflection achievable with one hexapole unit, a 12-segment NdFeB Halbach array with remanence $B_0 = 1.3\,\mathrm{T}$, internal radius $r_i = 3\,\mathrm{mm}$, external radius $r_e = 6\,\mathrm{mm}$ and thickness $l_h = 7\,\mathrm{mm}$, is characterised. The array is mounted on a linear actuator coupled with a z translation stage (to allow for linear movement parallel to the y axis as well as rotation in the yz plane), and positioned 6 mm beyond the last solenoid coil of the decelerator (Figure 4.2(a)). The Zeeman decelerator is operated using a standard decelerating sequence (with $\phi_0 = 90°$) that slows down hydrogen atoms to a target final velocity of $205\,\mathrm{ms}^{-1}$, generating the ToF trace shown in Figure 4.3.

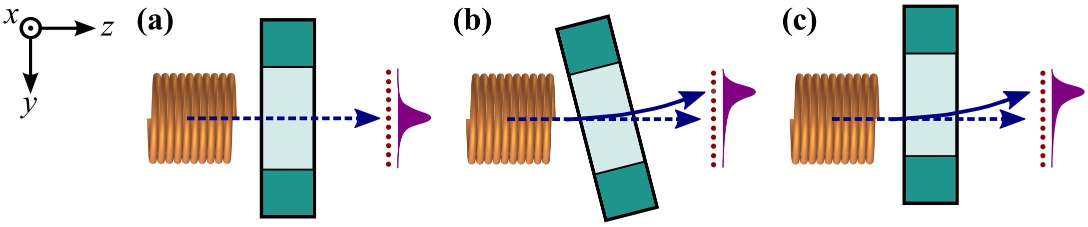

Figure 4.2: Schematic representation of the relative position of the Halbach array (in green, side view) when it is perfectly aligned (a), tilted (b) and shifted (c) with respect to the final decelerator coil (in orange). The blue arrows represent the trajectories of the deflected (solid line) and undeflected (dashed line) beams. The vertical red dots represent the detection positions probed in the characterization experiments for one Halbach array and the schematic beam profiles (in purple) represent the undeflected (a) and deflected (b and c) intensity profiles obtained.

The transmittance and deflection of particles as a function of the amount of tilt (Figure 4.2(b)) or shift (Figure 4.2(c)) of the Halbach array with respect to the last decelerator coil is investigated by detecting H atoms at a range of different positions along the y axis (see Figure 2.3(b)) at $z = 263\,\text{mm}$ (where $z = 0\,\text{mm}$ is the excimer laser position). This enables one to build up a profile of the spatial distribution of H atoms after the Halbach array (purple profiles in Figure 4.2). H atoms are detected by moving the ionisation laser vertically from $y = -5\,\text{mm}$ to $y = +5\,\text{mm}$ (with $y = 0\,\text{mm}$ corresponding to the decelerator axis) in steps of 1 mm, as represented by the red dots in Figure 4.2. At each detection position, the intensity of the signal from the decelerated peak is recorded at 10 µs intervals over a range of 50 µs, as represented by the vertical dashed lines in Figure 4.3. The integrated signal (i.e. the sum of the six measurements) for each scan is subsequently plotted as a function of detection position along the y axis, and a Gaussian profile is fitted through the experimentally measured signal levels. This process is repeated for a range of tilting angles (0.0°, ±6.25°, ±12.5°) and shifting distances (0.0 mm, ±0.2 mm, ±0.5 mm, ±1.0 mm), where a tilt of 0.0° and a shift of 0.0 mm correspond to the Halbach array being perfectly aligned with the last decelerator coil (Figure 4.2(a)). The maximum of each Gaussian and its uncertainty are shown as dots with respective error bars in Figure 4.4, where displacement of the decelerated peak in the y direction is plotted as a function of the amount of tilt (red colour) and shift (blue colour). The solid lines represent the predictions from three-dimensional particle-trajectory simulations, in which four million

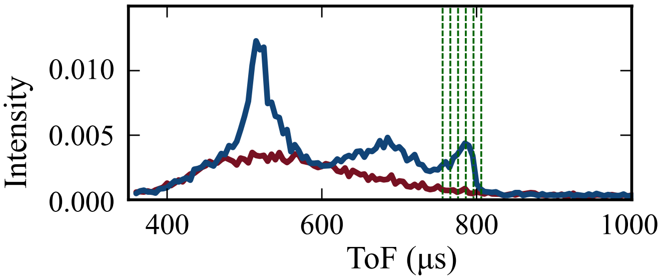

Figure 4.3: Experimental Time-of-Flight (ToF) traces for a beam of hydrogen atoms detected at $z = 263\,\mathrm{mm}$ when no magnetic field is applied (red) and when magnetic fields are applied to decelerate a subset of the beam to a target final velocity of $205\,\mathrm{m\,s^{-1}}$ (blue). The vertical dashed lines represent the scanning of the fully decelerated peak in order to record the intensity of the signal as a function of y detection position.

particles emerging from the decelerator are flown through a tilted or shifted Halbach array (see section 2.3). For each y detection position, the simulated signal resulting from the particles arriving at the detection region within the ToF window experimentally probed is integrated. Analogously to the experimental analysis, a Gaussian is fitted through the simulated signal levels as a function of detection position along the y axis, and the results are displayed in Figure 4.4. Tilting angles are simulated between $\pm 15^\circ$ in steps of 1°, whilst shifting distances are simulated between $\pm 1.2\,\mathrm{mm}$ in steps of $0.05\,\mathrm{mm}$.

Increasing the absolute shifting distance of the Halbach array with respect to the beam axis ($y = 0\,\mathrm{mm}$) displaces the decelerated peak along y in a linear fashion, and this behaviour is well reproduced by the simulations (Figure 4.4). The displacement of the decelerated peak as a function of the absolute tilting angle of the Halbach array with respect to the last decelerator coil is not captured by the simulations as well as the shifting behaviour. This is attributed to the broader beam profiles obtained from tilting (both in experiments and simulations) compared to shifting of the array, which render the Gaussian fitting less reliable. Indeed, when fitting a Gaussian through a set of measured intensities, a given fluctuation in one of the data points will have a larger effect on the fitting of a broad peak compared to that of a narrow peak, since the difference in intensity between neighbouring data points is smaller in the broad peak. However, in both experiment and simulation, the maximum displacement achievable with tilting is less than that achievable with shifting, within the accessible tilting and shifting ranges. Therefore, shifting deflects

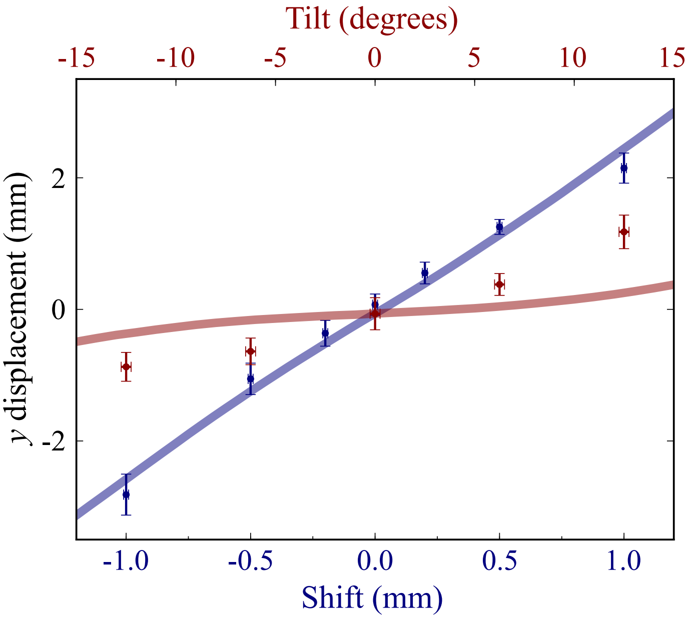

Figure 4.4: Displacement of decelerated-peak particles in the y direction as a function of tilt (red, top axes) and shift (blue, bottom axes) of one Halbach array (detected at $z = 263\,\mathrm{mm}$). The dots represent the experimental measurements (with error bars indicating the respective uncertainties) and the lines correspond to the simulated data (with the line width representing the uncertainty).

the decelerated peak more, as well as dispersing it less than tilting does. For these reasons, shifting of the Halbach arrays was chosen as the optimal deflection method in the design of the magnetic guide.

4.3 Design of the magnetic guide

Decelerated particles can be deflected using a single Halbach array shifted off axis, as demonstrated in section 4.2. In order to focus and return these decelerated particles to the initial beam axis, several Halbach arrays can be arranged in series. The arrays have a velocity-dependent focal length that is controlled by magnetisation, internal radius and thickness – similar to a wavelength-dependent optical focal length. Simulations evaluating the effect of remanence B_0, internal radius r_i, external radius r_e and thickness l_h of the arrays on focal length indicate that Halbach arrays with $B_0 = 1.3\,\mathrm{T}$, $r_i = 4\,\mathrm{mm}$, $r_e = 7\,\mathrm{mm}$ and $l_h = 5\,\mathrm{mm}$ are suitable to target H atoms with final velocities between

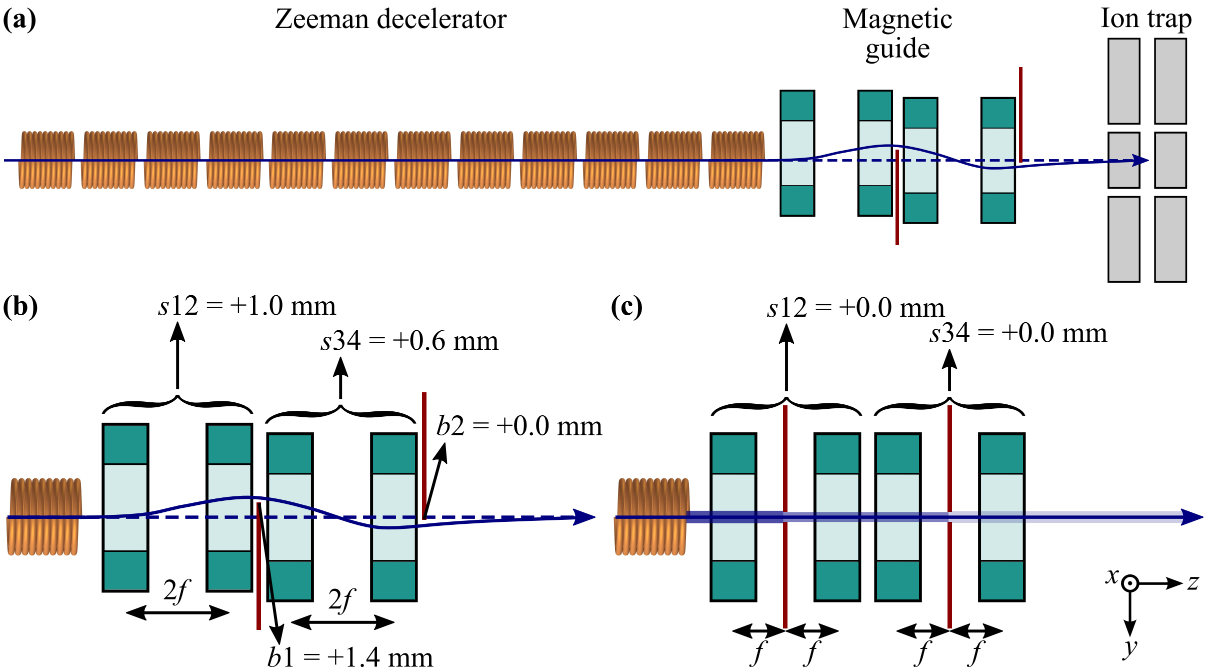

Figure 4.5: (a) Schematic representation of the Zeeman decelerator combined with the magnetic guide and an ion trap (not to scale). The red vertical lines represent the retractable blades whilst the blue arrows represent the trajectories of the deflected (solid line) and undeflected (dashed line) particles. (b) Geometry of the shifted guide proposed here, and (c) of its linear equivalent. The labels $s12$ and $s34$ correspond to the vertical shift of the first two Halbach arrays and of the second two arrays, respectively. In (b), $b1$ is the y position of the top end of the first blade and $b2$ is the y position of the bottom end of the second blade. The distance between the midpoint of arrays 1 and 2, and that of arrays 3 and 4, is equal to $2f$, where f is the focal length of the Halbach array for the velocity of interest. The blue arrows represent the trajectories of the deflected (solid line) and undeflected (dashed line) particles. In (c), the distance between the midpoint of each array and the nearest pinhole (in red) is equal to f. The blue arrow represents the trajectories of all particles whilst the superimposed shaded lines represent the gradual decrease in the number of faster-moving particles.

125 and 200 ms^{-1} – the target range for the Oxford Zeeman decelerator. Hence, Halbach arrays with these properties are chosen for the guide.

The optimal magnetic guide, determined using trajectory simulations with the aim of optimising the beam reaching the anticipated ion trap entrance at $z = 350$ mm, is composed of four Halbach arrays and two skimming blades, as depicted in Figure 4.5(a) and (b). The first array from the left deflects the slow particles up, the second collimates the deflected beam, the third brings it back to the beam axis and the fourth array collimates the beam a final time. The distance between the midpoints of the first two and of the second

Target velocity (ms^{-1})	125	150	175	200
Focal length, f (mm)	5.0	7.2	9.6	12.4

Table 4.1: Simulated focal lengths corresponding to a series of target velocities in the operational range of the guide. These are determined by simulating the trajectories (see section 2.3) of ten particles flying through one Halbach array with longitudinal velocity equal to the target velocity, and no transverse velocity. The particles are equally spaced in the transverse direction between $\pm 1\,\mathrm{mm}$ with respect to the beam axis.

two arrays is twice the focal length, $2f$ (see Table 4.1). Within each pair, the two arrays have the same vertical shift, which is larger for the first pair ($s12 = +1.0\,\mathrm{mm}$) than for the second pair ($s34 = +0.6\,\mathrm{mm}$) (see Figure 4.5(b)). There are two blades, placed after the second and fourth arrays, that serve to physically block the undeflected components of the beam. The blades penetrate beyond the central axis of the decelerator (the end of the first blade is at $y = b1 = +1.4\,\mathrm{mm}$ and the second blade extends down to the beam axis $y = b2 = +0.0\,\mathrm{mm}$), such that they overlap along the y axis leaving no direct line of sight between the decelerator and the ion trap. Hence, as demonstrated below, the blades can fully skim off all radical species that are not travelling at the target velocity, in addition to any seed gas and precursor molecules that reach the end of the decelerator. Only the radicals with a specific target velocity, corresponding to the specific focal length of the guide, follow the correct trajectory and pass through the guide unhindered.

Different target velocities can be addressed by altering the separation between the first and second arrays and between the third and fourth arrays, to account for the focal length variation with velocity (see Table 4.1). This is achieved by translating the second and fourth arrays along the z axis. Additionally, the blades can be retracted or extended to tune the purity of the emerging beam: increasing the overlap of the two blades improves the velocity selectivity of the guide, but also decreases the number of transmitted particles. The compromise between velocity resolution and number of transmitted particles is also crucial in the choice of the vertical shifts of the Halbach arrays: a larger vertical shift of the first pair of arrays leads to a narrower velocity distribution of the transmitted beam, at the expense of a lower number of transmitted particles.

The various degrees of freedom in the magnetic guide – the positions of Halbach arrays two and four (with respect to arrays one and three), the y-axis shift of the two pairs of arrays, and the y-axis penetration of the retractable blades – are controlled by vacuum-compatible, micrometer-adjustable actuators. Four linear actuators independently shift each pair of arrays and each blade in the y dimension whilst two rotary actuators turn two bevelled gears that are coupled to a threaded spindle to move the second and fourth arrays in the z dimension. In this way, the parameters of the guide can be adjusted whilst the apparatus remains under vacuum.

4.4 Results

4.4.1 Trajectory simulations

The velocity filtering of a beam of hydrogen atoms passing through the guide after deceleration – using an evolutionary-algorithm-optimised sequence – to a final velocity of $200\pm10\,\mathrm{ms}^{-1}$ ($f = 12.4\,\mathrm{mm}$) is illustrated in Figure 4.6 using trajectory simulations. Particles with the target velocity (top panel) are deflected up and then down by the arrays; they avoid both skimming blades and reach the target (detection) region where the ion trap entrance is anticipated to be (grey line at $z = 350\,\mathrm{mm}$). In contrast, faster-moving H atoms are completely skimmed off by the blades (middle panel) and slower-moving H atoms are partially skimmed off and partially deflected away from the target region (bottom panel).

The simulated ToF distribution of the particles reaching the detection region when the guide is present after the decelerator, compared to when it is not, is shown in Figure 4.7 (left panel). In the absence of the guide, the ToF trace has three peaks, as experimentally observed in Figure 4.3. When the guide is present, however, only part of the fully decelerated peak is transmitted, corresponding to particles with a narrow velocity distribution of $205\pm9\,\mathrm{ms}^{-1}$ (Figure 4.7, right panel). This guided peak arrives at slightly longer ToF

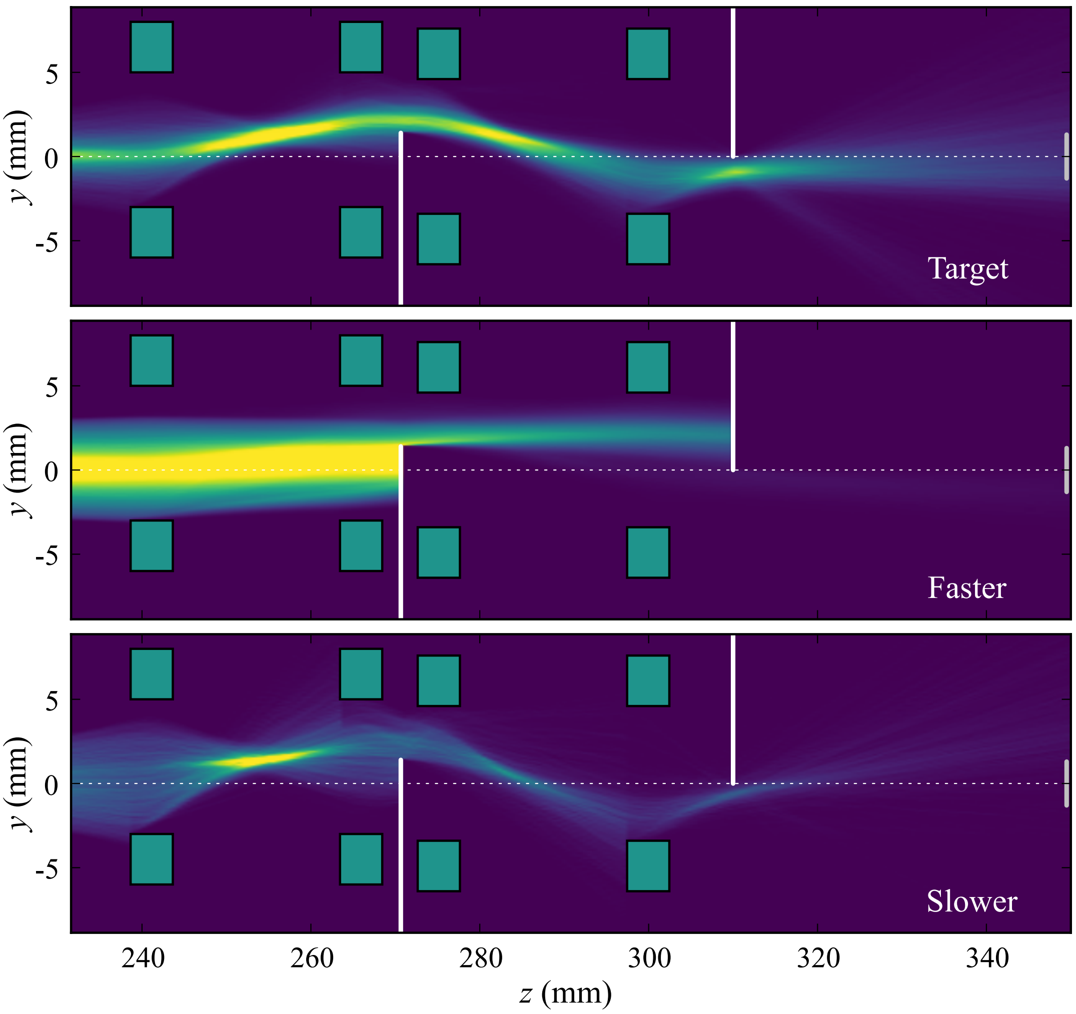

Figure 4.6: Three-dimensional numerical particle-trajectory simulations of a beam of 1s $^2S_{1/2}$ hydrogen atoms flying through the guide. The incoming beam has already passed through a 12-stage Zeeman decelerator with a target final velocity of $200\pm10\,\mathrm{ms^{-1}}$, and can be divided into three contributions: particles travelling at the target velocity, particles travelling faster than the target velocity, and particles travelling slower than the target velocity. The trajectories of the target particles ($v_z = 200\pm10\,\mathrm{ms^{-1}}$, top panel), faster particles ($v_z > 210\,\mathrm{ms^{-1}}$, middle panel) and slower particles ($v_z < 190\,\mathrm{ms^{-1}}$, bottom panel) through the magnetic guide are shown, with intensity ranging from low (purple) to high (yellow). The green rectangles represent the Halbach arrays in cross section, the white solid lines represent the blades and the grey line at $z = 350\,\mathrm{mm}$ represents the detection region. The dotted horizontal white line shows the decelerator axis (z axis) on which the beam is initially centred. Only the hydrogen atoms travelling within the narrow target velocity range are directed into the detection region.

compared to the corresponding peak when no guide is present (Figure 4.7, insets) because the guide induces a small decrease in the z-axis velocity of the decelerated particles. This

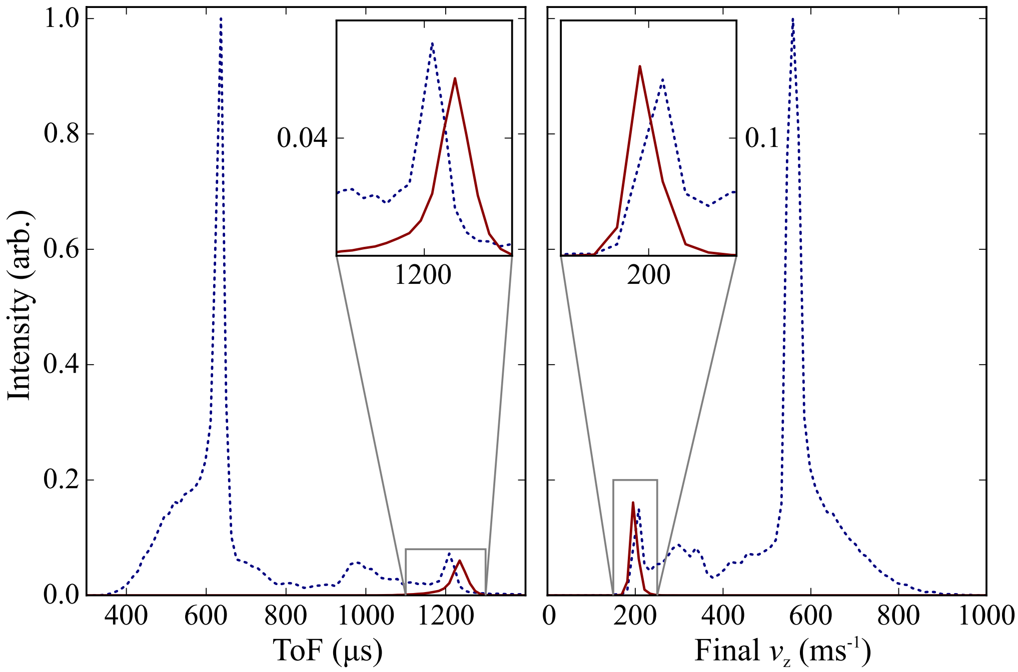

Figure 4.7: Simulated ToF (left panel) and velocity (right panel) distributions of H atoms arriving at the detection region ($z = 350\,\mathrm{mm}$) after passing through the guide (red solid line) compared to when no guide is present (blue dashed line).

is a direct result of the deflection in the y direction induced by the shifted arrays: when the particles are deflected away from the beam axis, their radial velocity increases at the expense of their axial velocity (Figure 4.8). For instance, a particle travelling at $y = 1.0\,\mathrm{mm}$ through the centre of the first array (which is shifted in y by $s12 = +1.0\,\mathrm{mm}$) is not deflected and its radial and axial velocities are constant (Figure 4.8, dark blue trace). On the other hand, a particle travelling at $y = -0.5\,\mathrm{mm}$ through the first array experiences a force towards the centre of the array to minimise its Zeeman energy such that its radial velocity is increased by $25\,\mathrm{ms}^{-1}$ – whilst its axial velocity is simultaneously decreased by $2\,\mathrm{ms}^{-1}$ – and the particle is deflected upwards (Figure 4.8, light blue trace).

Figure 4.8 shows the simulated trajectories of three hypothetical H atoms with the same velocity but different initial y position travelling through the guide. The particle with the largest displacement with respect to the centre of the first Halbach array ($y = -0.5\,\mathrm{mm}$, light blue) experiences the largest magnetic field from all four arrays (panel (e)) and undergoes the largest y displacement when passing through the guide (panel (a)). The speed

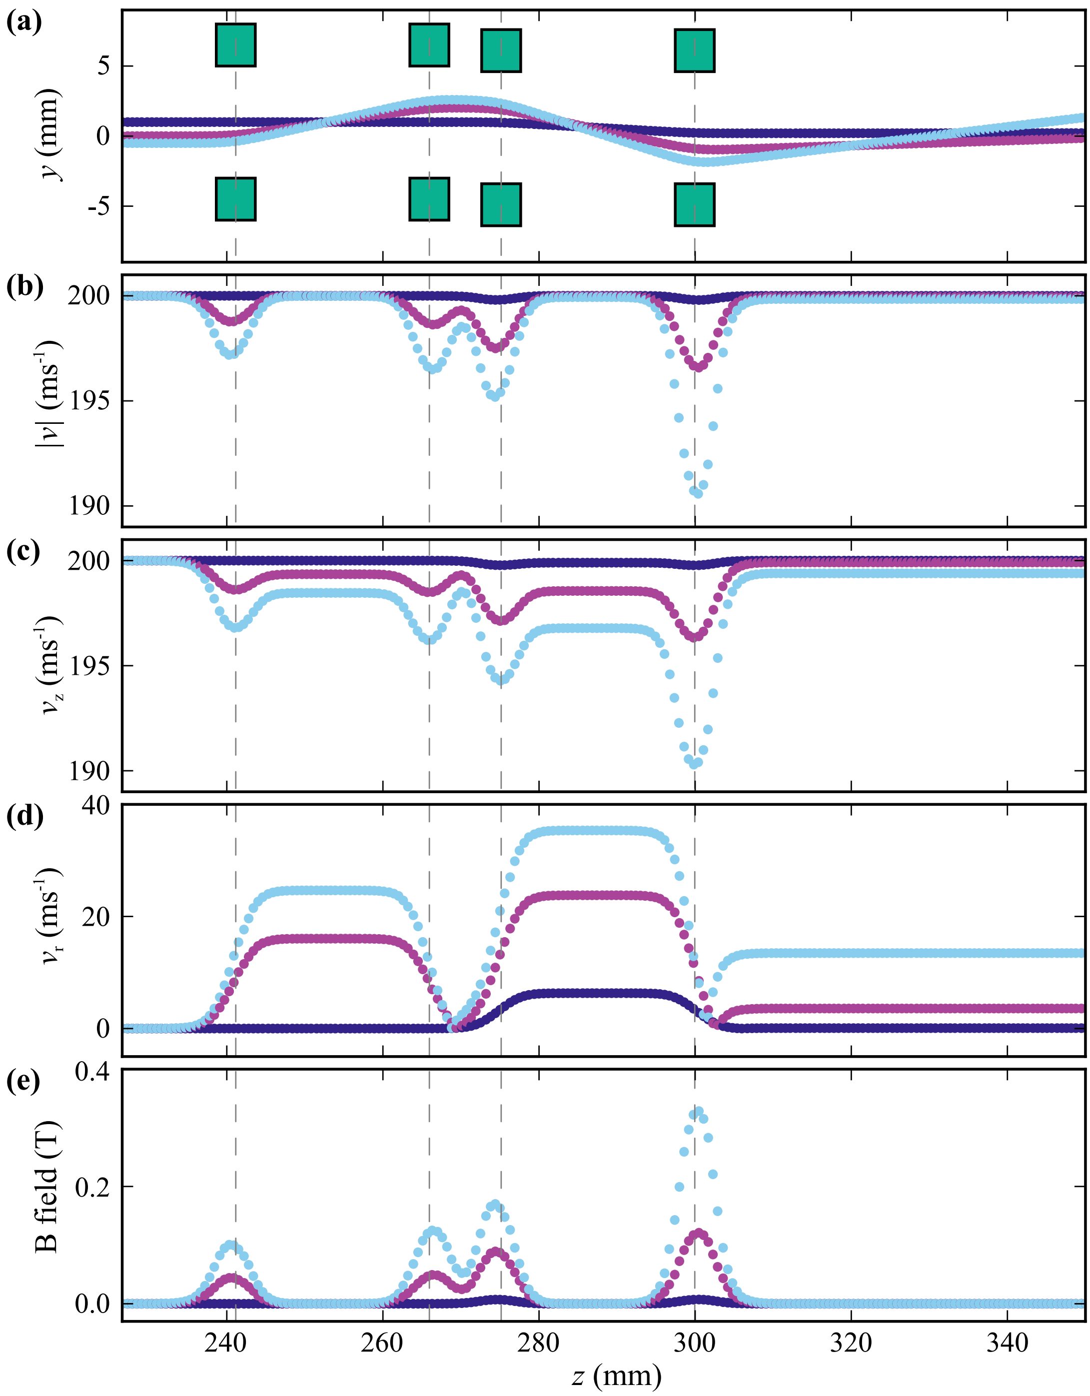

Figure 4.8: (a) Simulated trajectories of three hypothetical hydrogen atoms travelling through the guide, with the green rectangles representing the Halbach arrays in cross section. The particles are initially located at $x = 0.0\,\mathrm{mm}$ (all) and $y = -0.5\,\mathrm{mm}$ (light blue), $y = 0.0\,\mathrm{mm}$ (purple), $y = 1.0\,\mathrm{mm}$ (dark blue) with velocity $v_x = v_y = 0\,\mathrm{ms}^{-1}$, $v_z = 200\,\mathrm{ms}^{-1}$ (all). The speed (b), axial (c) and radial (d) velocities of the particles, as well as the magnetic field they experience (e), are also plotted as a function of z position. The dashed vertical lines indicate the centre of each array.

of the particle is reduced every time it travels towards the centre of an array and is increased again – by the same amount – as it travels away from the centre of the array (panel (b)). However, some of its axial velocity (panel (c)) gets converted into radial velocity (panel (d)), causing the displacement from the beam axis observed in panel (a). The same happens, to a lesser extent, to the particles starting at $y = 0.0\,\mathrm{mm}$ (purple) and $y = 1.0\,\mathrm{mm}$ (dark blue). As a result, the average radial velocity of the H atoms reaching the ion trap (at $z = 350\,\mathrm{mm}$) after passing through the guide is simulated to be $7.1\,\mathrm{ms}^{-1}$, in contrast to an average radial velocity of $2.6\,\mathrm{ms}^{-1}$ in the absence of the guide.

4.4.2 Performance of the magnetic guide

Simulations reveal that, of the hydrogen atoms that reach the detection region ($z = 350\,\mathrm{mm}$) following deceleration when the guide is not present, only 1.9% of the beam is within the target velocity range ($v_z = 200 \pm 10\,\mathrm{ms}^{-1}$); 98.0% of the beam is composed of faster particles and 0.1% of slower particles. After passing through the magnetic guide, 99% of the hydrogen atoms travelling outside the target velocity range are removed. Additionally, owing to the focusing properties of the Halbach arrays, the number of target particles that are transmitted through to the detection region is increased by 62% compared to the number of target particles that reach the detection region with no guide in place. It is geometrically not possible for the non-magnetic species (precursor and seed gases) to pass through the guide unhindered by the blades and, therefore, 100% of these species are removed from the beam. In practice, a differential-pumping stage is needed between the guide and the ion trap to prevent the skimmed particles from reaching the target region after bouncing off the blades.

When different velocities are guided – by varying the target velocity of the optimised deceleration sequence and the focal length of the guide – the simulated percentage of H atoms with target, faster and slower velocities is consistently around (81±3)%, (17±3)% and (2±1)% (respectively) across the 125–200 ms^{-1} operational range of the guide (Figure 4.9(a)). The average velocity of the beam emerging from the guide is in agreement

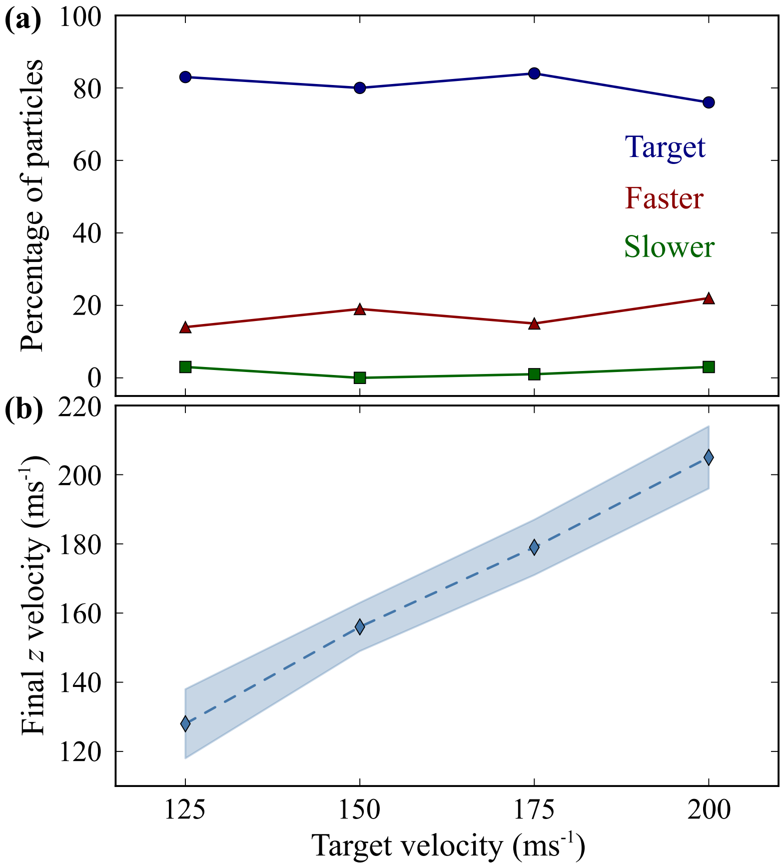

Figure 4.9: (a) Simulated percentage of particles with target (blue circles), faster (red triangles) and slower (green squares) velocities emerging from the guide (and detected at $z = 350\,\mathrm{mm}$) for each target velocity of deceleration and guiding. (b) Average velocity of the particles transmitted through the guide for each target velocity (light blue diamonds), with the shaded area representing one standard deviation.

with the target velocity (the expected beam velocity given the focal length of the guide) and, crucially, the velocity distribution of the beam is very narrow across the target velocity range (Figure 4.9(b)).

A linear equivalent of the magnetic guide described above could be implemented (Figure 4.5(c)), with the four Halbach arrays collinear and centred on the decelerator axis ($s12 = s34 = +0.0\,\mathrm{mm}$). In place of blades, pinholes could be placed at the focal points of the target velocity (that is, between the first and second arrays and between the third and fourth arrays), to block out the unwanted components of the beam. Such a linear design, while perhaps more intuitive than the approach adopted in this work, is significantly less efficient at purifying the beam. Even with a very small pinhole diameter of

0.5 mm, the percentage of H atoms with the target velocity after the guide is simulated to be only 20.9%, as the resulting beam is still mainly composed of faster-travelling H atoms (75.4%) that remain along the beamline axis. Furthermore, these values do not include the contribution to the transmitted beam of any precursor or seed gases travelling along the centre of the apparatus. This is because the pinholes are (necessarily) centred on the beam axis; even with a very small diameter and when positioned at the focal point for the target velocity, pinholes admit a large number of unwanted particles, given that the latter are several orders of magnitude more numerous than the target particles.

The shifted-guide design proposed in this chapter (Figure 4.5(b)) outperforms its linear equivalent (Figure 4.5(c)) due to the breaking of the *xy* symmetry which permits to leave no direct line of sight through the guide (Figure 4.10). This is essential to skim off all of the unwanted particles and achieve a high purity of slow H atoms. A linear design only achieves a gradual purification of the beam and thus would require a large number of stages to achieve a comparable purification to that obtained by the shifted design employed in this work, and would ultimately transmit far fewer target-velocity particles.

The shifted design also has several advantages over the bent design previously proposed [109]. Bending the decelerated particles away from the beam axis would require one to change the radius of the bend every time a different velocity is desired or a different species is to be guided. One would also need to move the ion trap with the guide in order to study ion–radical collisions, which is impractical from an experimental point of view. The new design circumvents this practical difficulty by maintaining a collinear arrangement of the trap with respect to the decelerator, as well as allowing for the tuning of the target velocity *in situ*, without having to break the vacuum.

Another alternative to using a magnetic guide to manipulate the radicals' trajectories would be to use a mechanical chopper or shutter to select only a portion of the beam based on arrival time. Mechanical shutters are employed by several research groups in combination with, for example, a Stark decelerator to physically block out unwanted elements in the beam [112]. The limited lifetime of mechanical shutters operating in

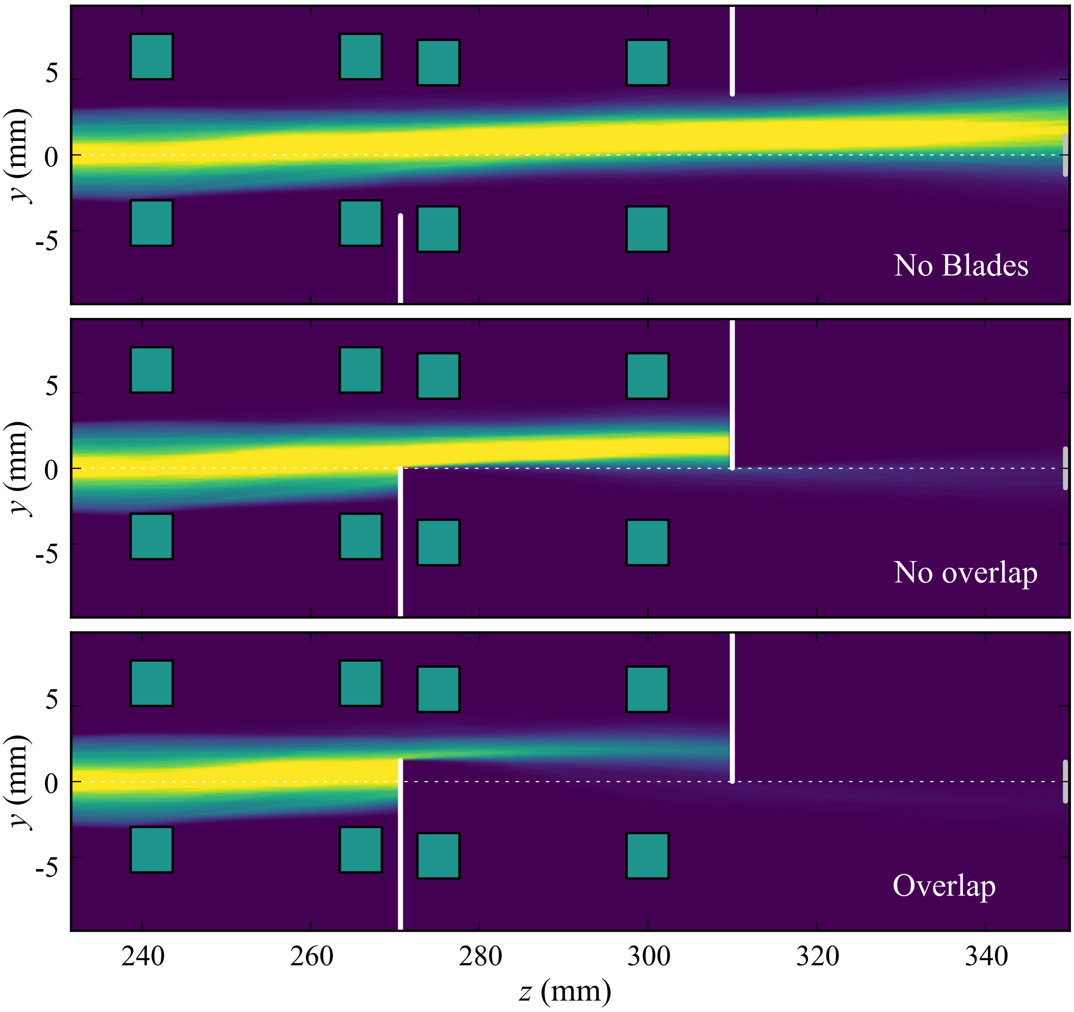

Figure 4.10: Three-dimensional numerical particle-trajectory simulations of a beam of 1s $^2\mathrm{S}_{1/2}$ hydrogen atoms flying through the guide after deceleration to the target final velocity of $200\pm10\,\mathrm{ms}^{-1}$. Only the trajectories of the particles travelling faster than the target velocity ($v_z > 210\,\mathrm{ms}^{-1}$) are shown flying through the guide: with the blades fully retracted ($b1 = -4.0\,\mathrm{mm}$, $b2 = +4.0\,\mathrm{mm}$; top panel, labelled "No blades"), with the blades on the beam axis ($b1 = b2 = +0.0\,\mathrm{mm}$; middle panel, labelled as "No overlap"), and with the blades at their optimised positions ($b1 = +1.4\,\mathrm{mm}$, $b2 = +0.0\,\mathrm{mm}$; bottom panel, labelled "Overlap"). The intensity ranges from low (purple) to high (yellow). The green rectangles represent the Halbach arrays in cross section, the white solid lines represent the blades and the grey line at $z = 350\,\mathrm{mm}$ represents the detection region. The dotted horizontal white line shows the decelerator axis (z axis) on which the beam is initially centred.

ultra-high-vacuum conditions, and the finite opening time of the shutter vanes (approximately 750 µs for fast-opening designs), limit the timeframe over which experiments can be recorded and the range of velocities that can be addressed. However, even a flawless

ToF-gating device would not give a sufficiently narrow velocity distribution for the apparatus described in this thesis, since the limited length of the Zeeman decelerator ($< 13\,\mathrm{cm}$) prevents full spatial separation of the particles in the beam according to their velocity (see Figure 3.7). As a result, the decelerated and undecelerated packets could not be effectively separated by a mechanical device unless it was placed some considerable distance after the final coil. A mechanical ToF-gating device would also transmit all species travelling at the same speed as the decelerated radicals of interest (such as slow-moving seed gas or precursor molecules). As such, mechanical shutters or choppers are less-than-ideal alternatives.

4.4.3 Experimental verification

Experimental ToF traces – recorded with and without the magnetic guide in place after the decelerator – are shown in Figure 4.11(a). The Zeeman decelerator is operated using an evolutionary-algorithm-optimised sequence with a target velocity of $200\pm10\,\mathrm{ms}^{-1}$, and the H atoms are detected at $z = 319\,\mathrm{mm}$. In the absence of the guide, the ToF trace has three peaks, arising from: H atoms that have not been decelerated, H atoms that have only been partially decelerated, and finally the fully decelerated H atoms (from shorter to longer time-of-flight, respectively). When the guide is placed after the decelerator, only the fully decelerated H atoms are transmitted, as predicted by the simulations (see section 4.4.1). Particles that contribute to the first two ToF peaks no longer reach the detection region. The experimental ToF traces are in excellent agreement with the simulated ToF profiles at all velocities.

The versatility of the magnetic guide presented here allows for the target velocity of H atoms to be tuned from $200\,\mathrm{ms}^{-1}$ down to $125\,\mathrm{ms}^{-1}$ by simply adjusting the position of the second and fourth Halbach arrays (achieved using rotary actuators, and thus without the need for breaking the vacuum) in combination with using a Zeeman decelerator pulse sequence optimised for the desired velocity. This is successfully demonstrated in Figure

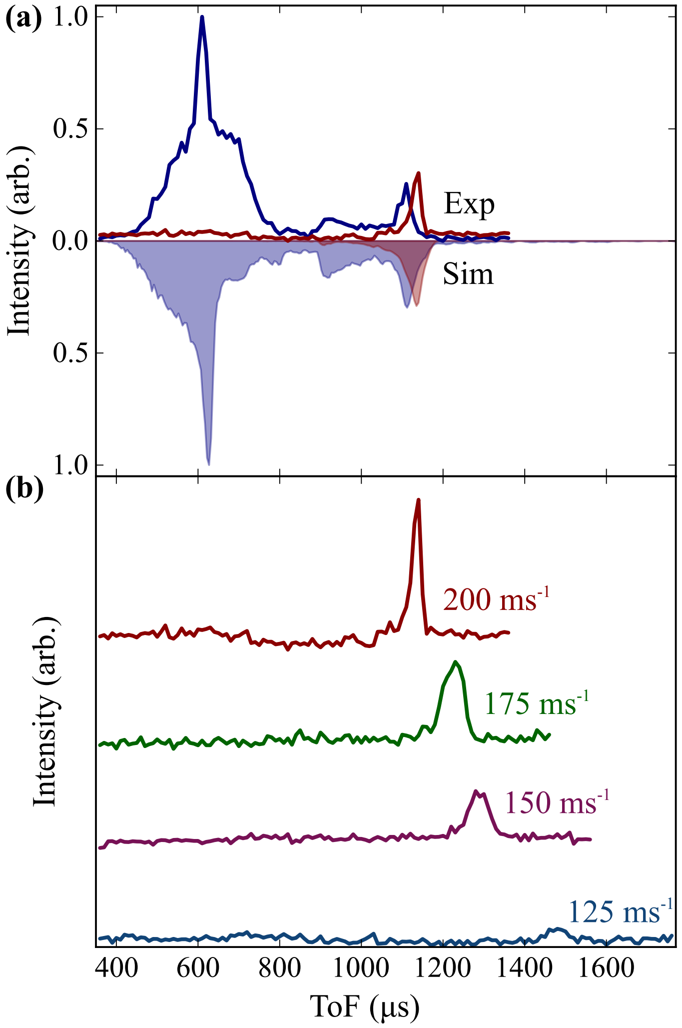

Figure 4.11: (a) Experimental ToF distribution of H atoms arriving at the detection region ($z = 319\,\mathrm{mm}$) after passing through the guide (red solid line) and when no guide is present (blue solid line), following deceleration to $200\,\mathrm{ms}^{-1}$. The corresponding simulated distributions are shown reflected below the experimental traces, with the red and blue shaded areas corresponding to the simulated ToF profiles with and without the guide present, respectively. (b) Experimental traces demonstrating that different final velocities can be targeted, spanning from $200\,\mathrm{ms}^{-1}$ down to $125\,\mathrm{ms}^{-1}$.

4.11(b), where the experimental ToF traces clearly exhibit a single peak that moves to later arrival times as the velocity of the H atoms that are transmitted is sequentially lowered.

The extent to which the blades penetrate the beam axis can also be modified to adjust the velocity resolution of the resulting guided beam: increasing the overlap between the blades increases the velocity resolution, at the expense of the number of particles transmitted (see section 4.3). This trade-off can be clearly seen in Figure 4.12: the trace

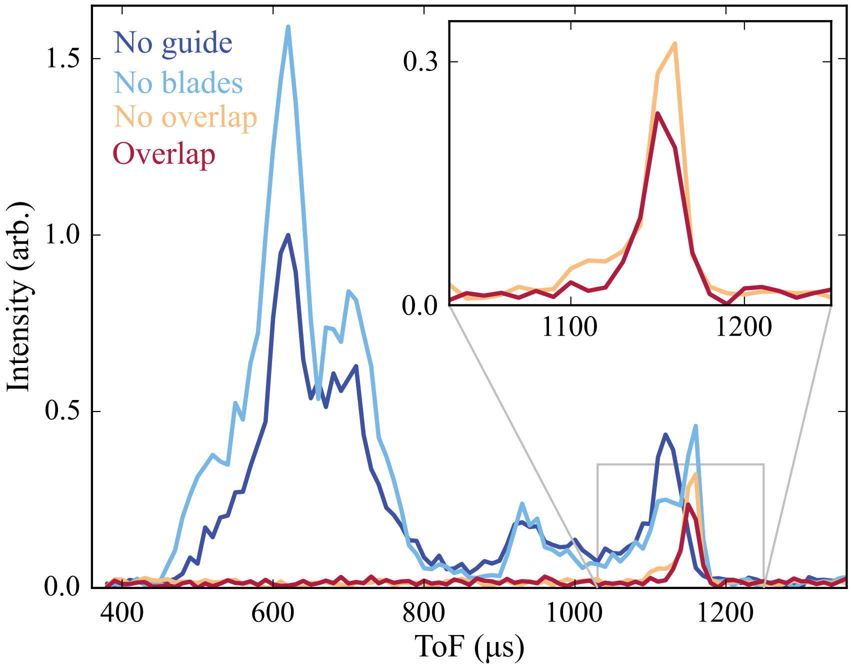

Figure 4.12: Experimental ToF distribution of H atoms arriving at the detection region ($z = 319\,\mathrm{mm}$) following deceleration to $200\,\mathrm{ms}^{-1}$ when no guide is present (dark blue trace, labelled "No guide") and after passing through the guide: with the blades fully retracted ($b1 = -4.0\,\mathrm{mm}$, $b2 = +4.0\,\mathrm{mm}$; light blue trace, labelled "No blades"), with the blades on the beam axis ($b1 = b2 = +0.0\,\mathrm{mm}$; orange trace, labelled as "No overlap"), and with the blades at their optimised positions ($b1 = +1.4\,\mathrm{mm}$, $b2 = +0.0\,\mathrm{mm}$; red trace, labelled "Overlap"). The inset shows only the "No overlap" and "Overlap" traces. (Note that the x axis does not match that in Figure 4.11 because a different pulsed-valve temperature of 213 K is used.)

recorded when the blades are fully retracted ($b1 = -4.0\,\mathrm{mm}$, $b2 = +4.0\,\mathrm{mm}$; light blue) is similar to the one where the guide is absent (dark blue). However, the former displays a larger signal due to the focusing effect of the arrays, and a shift to longer ToF of the fully decelerated peak due to the y deflection, as discussed in section 4.4.1. No velocity discrimination is achieved without blades. The trace recorded with the optimised blade overlap ($b1 = +1.4\,\mathrm{mm}$, $b2 = 0.0\,\mathrm{mm}$; red) is similar to the trace with no overlap of the blades (i.e. both blades positioned at precisely $b1 = b2 = +0.0\,\mathrm{mm}$; orange). The latter, however, exhibits a higher intensity of the decelerated peak and a higher-velocity shoulder (Figure 4.12, inset). The lower peak intensity and narrower arrival time (due to the disappearance of the higher-velocity shoulder) in the optimised configuration are attributed to the loss of partially decelerated H atoms. Indeed, simulations reveal that the average velocity of the hydrogen atoms reaching the ion trap (at $z = 350\,\mathrm{mm}$) when the blades are not overlapped is $275 \pm 136\,\mathrm{ms}^{-1}$, in contrast to an average velocity of $205 \pm 9\,\mathrm{ms}^{-1}$ when the blades are overlapped in the optimised configuration. The overlap of the blades is crucial for the achievement of the desired $\leq \pm 10\,\mathrm{ms}^{-1}$ velocity resolution.

Although the flux of H atoms emerging from the guide has not been experimentally determined yet, it is estimated to be comparable to the typical densities of decelerated particles obtained in conventional Zeeman decelerators, which are around $10^8\,\mathrm{cm}^{-3}$ [62]. Simulations indicate that about half of the particles with the target velocity emerging from the decelerator are successfully transmitted through the guide. Despite the relatively low number densities, the design of the magnetic guide is such that the user can decide upon (and modify) the optimal balance between transmission efficiency and velocity resolution for each given application *in situ*, based on the experimentally observed signal levels. This precise velocity filtering of radical beams is unprecedented, with the tuneability in the target velocity and resolution significantly increasing the potential functionality of the device.

Furthermore, a different velocity range (beyond the current capabilities) could be targeted – or a different paramagnetic species addressed – by exchanging the Halbach arrays described here for others with different radii (internal and external, Figure 4.1(a)), thickness and magnetisation. For instance, simulations indicate that Halbach arrays with $B_0 = 1.4\,\mathrm{T}$, $r_i = 3\,\mathrm{mm}$, $r_e = 7\,\mathrm{mm}$ and $l_h = 7\,\mathrm{mm}$ would be suitable to target H atoms with final velocities between 200 and $355\,\mathrm{ms}^{-1}$, as well as O atoms with final velocities between 85 and $155\,\mathrm{ms}^{-1}$. The simulations used for design optimisation are readily applicable to other species provided one has the relevant Zeeman energy level diagram for that species, and in general, the existing design for H atoms can be used as a starting point for optimising the geometry of the setup for other species. This opens the door to many interesting potential applications that would benefit from a pure source of radicals, with a tuneable velocity and narrow velocity distribution. Indeed, most radical sources yield a mixture of radical species and precursor gas and leave little scope for control over the velocity or purity of the beam.

While the magnetic guide was originally conceived to purify a Zeeman-decelerated beam of radicals, as successfully demonstrated here, it could also be employed as a stand-alone beam filter. Bent electrostatic quadrupole guides have long been employed to select those species moving with a velocity below a given threshold value from an effusive beam of dipolar particles [28, 113]. The magnetic guide presented here could do this – and more – by manipulating a collimated effusive beam of radicals. The latter could be generated, for instance, by photolysis of precursor molecules (possibly seeded in a carrier gas) within a capillary attached to a leak valve. Instead of transmitting all species moving below a certain threshold velocity, the flexible design of the magnetic guide will allow for narrow, and tuneable, velocity filtering of paramagnetic species from a collimated effusive source. In this way, a diverse range of radical beams could be generated, facilitating the study of radical reactions over a wide range of collision energies.

4.5 Conclusion

In summary, a magnetic guide composed of four z-translatable Halbach arrays and two y-translatable skimming blades has been demonstrated to act as a beam filter, such that only those radical species with a desired velocity are transmitted through the guide unhindered. Within the magnetic guide, the first two arrays are shifted in the y direction compared to the beam axis by +1.0 mm – to steer the H atoms travelling at the target velocity away from the beam axis and to collimate the deflected beam, respectively. The second pair of arrays is shifted in y by +0.6 mm, bringing the deflected radicals back to the beam axis and collimating the radical beam once again, respectively. Two retractable blades are positioned after the second and fourth arrays to remove all particles that do not follow the trajectory of the target H atoms. The breaking of the xy symmetry caused by the deflection away from (and back to) the beam axis allows the blades to overlap in the y dimension, yielding a considerably purer H-atom beam than can be achieved with a straight analogue of the proposed guide (Figure 4.5(c)). The target velocity is tuneable and can be varied by simply translating the position of the second and fourth Halbach arrays with respect to the first and third arrays. The properties of the arrays (internal and external radii, thickness and magnetisation) are chosen such that hydrogen atoms with velocities between $125\,\mathrm{ms}^{-1}$ and $200\,\mathrm{ms}^{-1}$ can be addressed. A different velocity range could be targeted, or a different radical species transmitted, by exchanging the current Halbach arrays for others with different radii, thickness or magnetisation.

Generating a pure source of velocity-selected radical species has been an ongoing experimental challenge, inhibiting the detailed experimental study of many important processes. The flexible magnetic-guide design detailed here achieves unprecedented purification of Zeeman-decelerated H atoms. The versatility of the magnetic guide is such that it could also be employed as a stand-alone device (without the Zeeman decelerator); provided there are some particles in the incoming beam that exhibit the desired velocity (as in an effusive beam), the magnetic guide will transmit only those particles that fall within the selected target velocity range. As such, the magnetic-guide design detailed here has a plethora of potential applications. With the single-ion sensitivity that can be achieved by monitoring reactions in Coulomb crystals, combining magnetic-guide-filtered radicals with Coulomb-crystallised ions held within an ion trap could pioneer the precise experimental examination of many important ion–radical reactions. Further applications, in areas spanning from fundamental measurements (chemical physics) to industrially-relevant reaction studies (radical-surface interactions), are envisaged to follow.

5 | Conclusions

The study of the reaction rate constants, temperature dependencies and product branching ratios of ion–radical reactions is vital to accurately model the chemistry of cold environments such as the interstellar medium. On the fundamental level, studying perfectly controlled reactions – with full knowledge of the velocities and quantum states of reactants and products – is the stepping stone to gaining full control over chemical reactions and manipulating their outcome using electromagnetic fields. Methods that are typically employed to investigate ion–radical reactions are afflicted by large uncertainties in the determination of reaction rate constants, partly due to the presence of precursor molecules in the radical beam. For instance, the rate constant for the reaction between H radicals and CO_2^+ ions was determined with an uncertainty of ±30% by SIFT [21] and of ±40% by AB-22PT [19]. A pure and controlled source of radicals is necessary to improve on these measurements, as well as to facilitate the detailed investigation of numerous other processes involving radicals.

Radicals are challenging species to purify since they are highly reactive, due to the presence of an unpaired electron in the outermost shell, and are formed in very low densities under typical laboratory conditions. As radical species do not possess a charge and frequently lack an appreciable electric dipole moment, they cannot be routinely controlled using electric fields. All radicals, however, possess a magnetic dipole moment and can therefore be manipulated using magnetic fields, to an extent determined by the mass and g-factor of the state in question. In this thesis, a source of pure, cold and state-selected radicals, composed of a Zeeman decelerator interfaced with a newly-designed magnetic guide, has been described.

Inside a Zeeman decelerator, some particles that are in low-field-seeking (LFS) quantum states are slowed down by a series of consecutive positive magnetic-field gradients, which convert part of the particles' kinetic energy into potential (Zeeman) energy that then gets dissipated. These magnetic fields are generated by passing current though a series of solenoid coils that compose the decelerator itself; they are spatially inhomogeneous – in order to create the required field gradients – and time-varying (i.e. pulsed) because they need to be switched off before the target particles enter the negative part of the field gradient. The series of timings for switching on and off the pulse of current applied to each of the solenoid coils compose a deceleration switching sequence.

Deceleration sequences are usually calculated by simulating the trajectory of one (synchronous) particle through the decelerator. In this thesis, evolutionary algorithms have been employed to optimise standard sequences and improve their performance. The evolutionary algorithm adopted here, the covariance matrix adaptation (CMA) evolutionary strategy (ES), starts from an input (standard) sequence and randomly modifies the durations of the pulses of current applied to each coil to generate a first population of trial sequences. Each trial sequence is then evaluated by simulating the trajectory of thousands of particles travelling thought the decelerator, and the number of particles that emerge with velocity within a certain target range is recorded. The worst-performing trial sequences are then discarded, whilst the best-performing sequences are used in the generation of the next population of sequences. This process is iterated until the number of particles decelerated to the target velocity range is maximized and convergence is reached, which is when the overall best performing sequence has not been updated for 50,000 generations. This type of optimization outperforms steepest-descent methods in high-dimensional search problems with a large number of coupled parameters such as this, and the convergence is reliably and rapidly reached by using information from the best-performing sequences in previous generations to guide the generation of new trial sequences.

The optimised deceleration sequences obtained are capable of reaching much lower velocities than can be achieved using standard sequences, under the same experimental conditions. Experimental measurements and simulations confirm that a beam of ground-state H atoms in the 1s ($^2\mathrm{S}_{1/2}$) state can be decelerated down to $75\,\mathrm{m s^{-1}}$ (corresponding to $T = 0.3\,\mathrm{K}$) with optimised sequences, when the lowest velocity attainable with standard sequences is $205\,\mathrm{m s^{-1}}$ (corresponding to $T = 2.5\,\mathrm{K}$). Additionally, for the same final velocity, the optimised sequences produce 40% more target particles than standard sequences. Three-dimensional particle-trajectory simulations reveal that this is due to the improved phase-space acceptance: particles with a wider range of initial velocities are targeted by the optimised sequences, which in this way compromise between targeting particles that are slower than the most probable velocity of the beam (to reach lower final velocities) and maximising the number of decelerated particles. The algorithm optimises sequences considering the properties of the entire beam and the passage of all particles through the decelerator, instead of only considering the trajectory of one synchronous particle. This leads to a two-stage deceleration mechanism in which faster particles are addressed first and slower particles are progressively added to the targeted bunch as it moves through the decelerator, rather than being part of the bunch from the beginning.

Future work is planned to understand whether the observed two-stage mechanism is only optimal for the specific deceleration of H atoms in a 12-stage decelerator, or whether it is observed consistently in sequences optimised for different-length decelerators and for different radical species. By studying the optimisation output for different decelerators and radical species, it might be possible to identify patterns that could eventually lead to engineering optimal sequences without the need for the optimisation algorithm. This would be advantageous especially in the case of long decelerators, where the computational time for evolutionary-algorithm optimisation could become a limiting factor. Another future goal consists in including additional optimisation criteria to simultaneously maximise the number of decelerated particles within a target range and, for example, minimise the particles with velocity close to but outside that range, or minimise the longitudinal spatial spread of the decelerated packet. Evolutionary-algorithm optimisation has the potential to offer a much greater degree of control over the properties of the decelerated beam than has so far been explored.

The first, proof-of-principle demonstration of Zeeman deceleration of metastable $\mathrm{N}(^2\mathrm{D}_{5/2})$ atoms from 460 to 410 $\mathrm{m\,s^{-1}}$ has been reported in this thesis. The rather small decrease in the particles' kinetic energy of 20% (compared to $> 90\%$ for H atoms) is due to the limited number of deceleration stages available. As a consequence, the decelerated packet is not observed as a separate peak (as in the deceleration of H atoms) but as a “shoulder” in a distribution that consists mainly of non-decelerated particles. Optimisation of the sequences for deceleration of metastable N atoms using evolutionary algorithms may allow more energy to be removed without requiring a longer decelerator.

In order to convert the beam emerging from the decelerator – which is composed of H atoms with a wide range of velocities, precursor molecules and seed gases – into a pure beam of LFS hydrogen atoms with a well-defined and narrow velocity distribution, a magnetic guide has been designed and constructed. Halbach arrays have been chosen as the deflecting elements and the deflection of slow particles in the beam has been thoroughly investigated. Shifting an array with respect to the beam axis has been shown to generate more deflection and, at the same time, less dispersion of the particles than tilting the array; shifting was therefore the chosen deflection method in the guide design.

The magnetic guide consists of two pairs of moveable Halbach arrays and two retractable skimming blades. All elements are adjustable *in situ*, without having to break the vacuum. The two pairs of arrays are shifted with respect to the beam axis such that particles travelling at the target velocity are deflected away from, and then back to, the beam axis. All other particles are not appropriately manipulated by the Halbach arrays and thus cannot traverse the guide unhindered; they are skimmed by the two blades. Adjusting the distance between the two pairs of arrays allows different velocities to be guided – within the range of 125 to 200 $\mathrm{ms^{-1}}$ (0.9 to 2.4 K) – whilst adjusting the relative positions of the retractable blades allows tuning of the velocity distribution of the transmitted particles.

Experimental measurements and trajectory simulations confirm that, for a target velocity of 200±10 $\mathrm{ms^{-1}}$, the magnetic guide removes 99% of the hydrogen atoms travelling outside the target velocity range, in addition to 100% of the non-magnetic particles. The transmitted beam has an average velocity of 205±9 $\mathrm{ms^{-1}}$ and 62% more particles reach the detection region compared to when no guide is present. Pure beams of H atoms with velocities down to 125 $\mathrm{ms^{-1}}$ are obtained experimentally by simply decreasing the distance between the arrays in each of the two pairs, and operating the decelerator using a sequence that is optimised for the desired velocity. Simulations reveal that a high purity and a narrow velocity distribution of transmitted particles around the target velocity are achieved across the operational range of the device.

Future developments include the extension of the current operational range of the guide to address H atoms with higher target velocities. This can be achieved by simply exchanging the current Halbach arrays for others with slightly different specifications, without having to modify the frame of the guide. For instance, simulations indicate that Halbach arrays with $B_0 = 1.4\,\mathrm{T}$, $r_i = 3\,\mathrm{mm}$, $r_e = 7\,\mathrm{mm}$ and $l_h = 7\,\mathrm{mm}$ are suitable to target H atoms with final velocities between 200 and $355\,\mathrm{ms}^{-1}$. Besides targeting H atoms that are travelling at higher final velocities, these new magnets will allow the magnetic guide to be used as a stand-alone velocity and quantum-state filter. Indeed, experiments are planned to show that H atoms can be velocity- and state-filtered directly from the undecelerated beam. Such a demonstration is not possible with the current Halbach arrays, as there are too few H atoms travelling at the upper limit of the velocity range that can be transmitted through the guide. Subsequently, the guide could be used on its own (without the Zeeman decelerator) to generate slow beams of state- and velocity-selected radicals other than hydrogen. For instance, the new magnets described above could also guide oxygen atoms with velocities between 85 and $155\,\mathrm{ms}^{-1}$ from a collimated effusive beam, which could be generated by photo-dissociating NO_2 within a capillary attached to a leak valve. Molecular radicals could potentially also be guided this way, although the distribution of rotational and vibrational states would be thermal if the particles are produced in an effusive, rather than supersonic, beam. Finally, it would be interesting to employ an evolutionary strategy to optimise the positions of the various elements in the magnetic guide, to see if its performance could be further improved. Simultaneous optimisation of the deceleration sequence and magnetic guide geometry could also be beneficial since it might shape the properties of the decelerator output to improve the efficiency of the guide.

In conclusion, the source of cold and state-selected radicals described in this thesis has been shown to generate a velocity-selected beam of ground-state H atoms that is free from precursor molecules and seed gas. The velocity tuneability of the source will allow studies of reaction rates as a function of collision energy to be performed, and the narrow velocity distribution of the generated beam will afford excellent collision-energy resolution. This source will shortly be combined with a cryogenic ion trap, where molecular ions such as CO_2^+ formed by REMPI can be sympathetically cooled into Ca^+ Coulomb crystals. The long trapping times and single-ion sensitivity of the latter, combined with the pure and controlled H-atom beam described in this thesis, will permit the detailed study of ion–radical reactions at low temperatures.

Bibliography

[1] Dulieu, O., Osterwalder, A., Eds. *Cold Chemistry: Molecular Scattering and Reactivity Near Absolute Zero*; Royal Society of Chemistry, 2018.

[2] Stuhl, B. K.; Hummon, M. T.; Ye, J. *Annu. Rev. Phys. Chem.* **2014**, *65*, 501.

[3] Bell, M. T.; Softley, T. P. *Mol. Phys.* **2009**, *107*, 99.

[4] Krems, R. V. *Phys. Chem. Chem. Phys.* **2008**, *10*, 4079.

[5] van de Meerakker, S. Y. T.; Bethlem, H. L.; Meijer, G. *Nat. Phys.* **2008**, *4*, 595.

[6] Ehrenfreund, P.; Charnley, S. B. *Annu. Rev. Astron. Astrophys.* **2000**, *38*, 427.

[7] Klemperer, W. *Annu. Rev. Phys. Chem.* **2011**, *62*, 173.

[8] Williams, D. A.; Brown, W. A.; Price, S. D.; Rawlings, J. M. C.; Viti, S. *Astron. Geophys.* **2007**, *48*, 25.

[9] Herbst, E. *Chem. Soc. Rev.* **2001**, *30*, 168.

[10] Herbst, E.; van Dishoeck, E. F. *Annu. Rev. Astron. Astrophys.* **2009**, *47*, 427.

[11] Williams, D.; Viti, S. *Philos. Trans. Royal Soc. A; Math., Phys. and Eng. Sci.* **2013**, *371*.

[12] Ehrenfreund, P.; Spaans, M.; Holm, N. G. *Philos. Trans. Royal Soc. A: Math., Phys. and Eng. Sci.* **2011**, *369*, 538.

[13] Anicich, V. G. *J. Phys. Chem. Ref. Data* **1993**, *22*, 1469.

[14] Snow, T. P.; Bierbaum, V. M. *Annu. Rev. Anal. Chem.* **2008**, *1*, 229.

[15] Sablier, M.; Rolando, C. *Mass Spectrom. Rev.* **1993**, *12*, 285.

[16] Smith, I. W. *Annu. Rev. Astron. Astrophys.* **2011**, *49*, 29.

[17] Wakelam, V.; Smith, I. W. M.; Herbst, E.; Troe, J.; Geppert, W.; Linnartz, H.; Öberg, K.; Roueff, E.; Agúndez, M.; Pernot, P.; Cuppen, H. M.; Loison, J. C.; Talbi, D. *Space Sci. Rev.* **2010**, *156*, 13.

[18] Adams, N.; Smith, D. *Int. J. Mass Spectrom. Ion Phys.* **1976**, *21*, 349.

[19] Borodi, G.; Luca, A.; Gerlich, D. *Int. J. Mass Spectrom.* **2009**, *280*, 218.

[20] Gerlich, D.; Borodi, G.; Luca, A.; Mogo, C.; Smith, M. A. *Z. Phys. Chem.* **2011**, *225*, 475.

[21] Scott, G. B.; Fairley, D. A.; Freeman, C. G.; McEwan, M. J.; Spanel, P.; Smith, D. *J. Chem. Phys.* **1997**, *106*, 3982.

[22] *Faraday Discuss.* **2009**, *142*, 93.

[23] Fournier, M.; Le Picard, S. D.; Sims, I. R. In *Cold Chemistry: Molecular Scattering and Reactivity Near Absolute Zero*; Dulieu, O., Osterwalder, A., Eds.; The Royal Society of Chemistry, 2018; Chapter 1, pp 1–45.

[24] Berteloite, C.; Lara, M.; Bergeat, A.; Le Picard, S. D.; Dayou, F.; Hickson, K. M.; Canosa, A.; Naulin, C.; Launay, J.-M.; Sims, I. R.; Costes, M. *Phys. Rev. Lett.* **2010**, *105*, 203201.

[25] Marquette, J.; Rowe, B.; Dupeyrat, G.; Poissant, G.; Rebrion, C. *Chem. Phys. Lett.* **1985**, *122*, 431.

[26] Rowe, B.; Canosa, A.; Le Page, V. *Int. J. Mass Spectrom. Ion Proc.* **1995**, *149-150*, 573 – 596, Honour Biography David Smith.

[27] Zhang, D.; Willitsch, S. In *Cold Chemistry: Molecular Scattering and Reactivity Near Absolute Zero*; Dulieu, O., Osterwalder, A., Eds.; The Royal Society of Chemistry, 2018; Chapter 10, pp 496–536.

[28] Willitsch, S.; Bell, M. T.; Gingell, A. D.; Procter, S. R.; Softley, T. P. *Phys. Rev. Lett.* **2008**, *100*, 043203.

[29] Heazlewood, B. R.; Softley, T. P. *Annu. Rev. Phys. Chem.* **2015**, *66*, 475.

[30] Hansen, A.; Versolato, O.; Kłosowski, Ł.; Kristensen, S.; Gingell, A.; Schwarz, M.; Windberger, A.; Ullrich, J.; López-Urrutia, J. C.; Drewsen, M. *Nature* **2014**, *508*, 76.

[31] Schmid, P. C.; Greenberg, J.; Miller, M. I.; Loeffler, K.; Lewandowski, H. J. *Rev. Sci. Instrum.* **2017**, *88*, 123107.

[32] Mølhave, K.; Drewsen, M. *Phys. Rev. A* **2000**, *62*, 011401.

[33] Roth, B.; Wenz, H.; Daerr, H.; Schiller, S. *Phys. Rev. A* **2006**, *73*, 042712.

[34] Okada, K.; Suganuma, T.; Furukawa, T.; Takayanagi, T.; Wada, M.; Schuessler, H. A. *Phys. Rev. A* **2013**, *87*, 043427.

[35] Willitsch, S.; Bell, M. T.; Gingell, A. D.; Softley, T. P. *Phys. Chem. Chem. Phys.* **2008**, *10*, 7200.

[36] Greenberg, J.; Schmid, P. C.; Miller, M.; Stanton, J. F.; Lewandowski, H. J. *Phys. Rev. A* **2018**, *98*, 032702.

[37] Meyer, K. A. E.; Pollum, L. L.; Petralia, L. S.; Tauschinsky, A.; Rennick, C. J.; Softley, T. P.; Heazlewood, B. R. *J. Phys. Chem. A* **2015**, *119*, 12449.

[38] Eberle, P.; Dörfler, A. D.; von Planta, C.; Ravi, K.; Haas, D.; Zhang, D.; van de Meerakker, S. Y. T.; Willitsch, S. *J. Phys.: Conf. Ser.* **2015**, *635*, 012012.

[39] Bethlem, H. L.; Berden, G.; Meijer, G. *Phys. Rev. Lett.* **1999**, *83*, 1558.

[40] Bochinski, J. R.; Hudson, E. R.; Lewandowski, H. J.; Meijer, G.; Ye, J. *Phys. Rev. Lett.* **2003**, *91*, 243001.

[41] van de Meerakker, S. Y. T.; Labazan, I.; Hoekstra, S.; Küpper, J.; Meijer, G. *J. Phys. B: At. Mol. Opt. Phys.* **2006**, *39*, S1077.

[42] Wang, X.; Kirste, M.; Meijer, G.; van de Meerakker, S. Y. T. *Z. Phys. Chem.* **2013**, *227*, 1595.

[43] Nourbakhsh, O.; Michan, J. M.; Mittertreiner, T.; Carty, D.; Wrede, E.; Djuricanin, P.; Momose, T. *Mol. Phys.* **2015**, *113*, 4007.

[44] van de Meerakker, S. Y. T.; Smeets, P. H. M.; Vanhaecke, N.; Jongma, R. T.; Meijer, G. *Phys. Rev. Lett.* **2005**, *94*, 023004.

[45] Hoekstra, S.; Metsälä, M.; Zieger, P. C.; Scharfenberg, L.; Gilijamse, J. J.; Meijer, G.; van de Meerakker, S. Y. T. *Phys. Rev. A* **2007**, *76*, 063408.

[46] Sawyer, B. C.; Lev, B. L.; Hudson, E. R.; Stuhl, B. K.; Lara, M.; Bohn, J. L.; Ye, J. *Phys. Rev. Lett.* **2007**, *98*, 253002.

[47] Willitsch, S. *Int. Rev. Phys. Chem.* **2012**, *31*, 175.

[48] Härter, A.; Hecker Denschlag, J. *Contemp. Phys.* **2014**, *55*, 33.

[49] Haze, S.; Hata, S.; Fujinaga, M.; Mukaiyama, T. *Phys. Rev. A* **2013**, *87*, 052715.

[50] Sivarajah, I.; Goodman, D. S.; Wells, J. E.; Narducci, F. A.; Smith, W. W. *Phys. Rev. A* **2012**, *86*, 063419.

[51] Hall, F. H. J.; Aymar, M.; Bouloufa-Maafa, N.; Dulieu, O.; Willitsch, S. *Phys. Rev. Lett.* **2011**, *107*, 243202.

[52] Eberle, P.; Dörfler, A. D.; von Planta, C.; Ravi, K.; Willitsch, S. *Chem. Phys. Chem.* **2016**, *17*, 3769.

[53] Ghaffari, B.; Gerton, J. M.; McAlexander, W. I.; Strecker, K. E.; Homan, D. M.; Hulet, R. G. *Phys. Rev. A* **1999**, *60*, 3878.

[54] Patterson, D.; Doyle, J. M. *J. Chem. Phys.* **2007**, *126*, 154307.

[55] Henson, A. B.; Gersten, S.; Shagam, Y.; Narevicius, J.; Narevicius, E. *Science* **2012**, *338*, 234.

[56] Jankunas, J.; Bertsche, B.; Jachymski, K.; Hapka, M.; Osterwalder, A. *J. Chem. Phys.* **2014**, *140*, 244302.

[57] Even, U. *Eur. Phys. J. Techn. Instrum.* **2015**, *2*, 17.

[58] Shagam, Y.; Narevicius, E. *J. Phys. Chem. C* **2013**, *117*, 22454.

[59] van der Poel, A. P. P.; Dulitz, K.; Softley, T. P.; Bethlem, H. L. *New J. Phys.* **2015**, *17*, 055012.

[60] van der Poel, A. P. P.; Zieger, P. C.; van de Meerakker, S. Y. T.; Loreau, J.; van der Avoird, A.; Bethlem, H. L. *Phys. Rev. Lett.* **2018**, *120*, 033402.

[61] Hogan, S. D.; Motsch, M.; Merkt, F. *Phys. Chem. Chem. Phys.* **2011**, *13*, 18705.

[62] Narevicius, E.; Raizen, M. G. *Chem. Rev.* **2012**, *112*, 4879.

[63] Narevicius, E.; Parthey, C. G.; Libson, A.; Narevicius, J.; Chavez, I.; Even, U.; Raizen, M. G. *New J. Phys.* **2007**, *9*, 358.

[64] Vanhaecke, N.; Meier, U.; Andrist, M.; Meier, B. H.; Merkt, F. *Phys. Rev. A* **2007**, *75*, 031402.

[65] Hogan, S. D.; Wiederkehr, A. W.; Schmutz, H.; Merkt, F. *Phys. Rev. Lett.* **2008**, *101*, 143001.

[66] Hogan, S. D.; Sprecher, D.; Andrist, M.; Vanhaecke, N.; Merkt, F. *Phys. Rev. A* **2007**, *76*, 023412.

[67] Narevicius, E.; Libson, A.; Parthey, C. G.; Chavez, I.; Narevicius, J.; Even, U.; Raizen, M. G. *Phys. Rev. A* **2008**, *77*, 051401.

[68] Momose, T.; Liu, Y.; Zhou, S.; Djuricanin, P.; Carty, D. *Phys. Chem. Chem. Phys.* **2013**, *15*, 1772.

[69] Akerman, N.; Karpov, M.; Segev, Y.; Bibelnik, N.; Narevicius, J.; Narevicius, E. *Phys. Rev. Lett.* **2017**, *119*, 073204.

[70] Cremers, T.; Chefdeville, S.; Plomp, V.; Janssen, N.; Sweers, E.; van de Meerakker, S. Y. T. *Phys. Rev. A* **2018**, *98*, 033406.

[71] Trimeche, A.; Bera, M. N.; Cromières, J.-P.; Robert, J.; Vanhaecke, N. *Eur. Phys. J. D* **2011**, *65*, 263.

[72] Motsch, M.; Jansen, P.; Agner, J. A.; Schmutz, H.; Merkt, F. *Phys. Rev. A* **2014**, *89*, 043420.

[73] Dulitz, K.; Tauschinsky, A.; Softley, T. P. *New J. Phys.* **2015**, *17*, 035005.

[74] Dulitz, K.; Toscano, J.; Tauschinsky, A.; Softley, T. P. *J. Phys. B: At., Mol. Opt. Phys.* **2016**, *49*, 075203.

[75] Lavert-Ofir, E.; Gersten, S.; Henson, A. B.; Shani, I.; David, L.; Narevicius, J.; Narevicius, E. *New J. Phys.* **2011**, *13*, 103030.

[76] Cremers, T.; Chefdeville, S.; Janssen, N.; Sweers, E.; Koot, S.; Claus, P.; van de Meerakker, S. Y. T. *Phys. Rev. A* **2017**, *95*, 043415.

[77] Akerman, N.; Karpov, M.; David, L.; Lavert-Ofir, E.; Narevicius, J.; Narevicius, E. *New J. Phys.* **2015**, *17*, 065015.

[78] Toscano, J.; Tauschinsky, A.; Dulitz, K.; Rennick, C. J.; Heazlewood, B. R.; Softley, T. P. *New J. Phys.* **2017**, *19*, 083016.

[79] Toscano, J.; Rennick, C. J.; Softley, T. P.; Heazlewood, B. R. *J. Chem. Phys.* **2018**, *149*, 174201.

[80] Toscano, J.; Hejduk, M.; McGhee, H. G.; Heazlewood, B. R. *In preparation.*

[81] Dulitz, K. Towards the study of cold chemical reactions using Zeeman decelerated supersonic beams. Ph.D. thesis, University of Oxford, 2014.

[82] Wiederkehr, A. W.; Hogan, S. D.; Merkt, F. *Phys. Rev. A* **2010**, *82*, 043428.

[83] van de Meerakker, S. Y. T.; Vanhaecke, N.; Bethlem, H. L.; Meijer, G. *Phys. Rev. A* **2006**, *73*, 023401.

[84] Dulitz, K.; Vanhaecke, N.; Softley, T. P. *Phys. Rev. A* **2015**, *91*, 013409.

[85] Zhang, D.; Meijer, G.; Vanhaecke, N. *Phys. Rev. A* **2016**, *93*, 023408.

[86] Dulitz, K.; Motsch, M.; Vanhaecke, N.; Softley, T. P. *J. Chem. Phys.* **2014**, *140*, 104201.

[87] Wiederkehr, A. W.; Motsch, M.; Hogan, S. D.; Andrist, M.; Schmutz, H.; Lambillotte, B.; Agner, J. A.; Merkt, F. *J. Chem. Phys.* **2011**, *135*, 214202.

[88] Bergeman, T.; Erez, G.; Metcalf, H. J. *Phys. Rev. A* **1987**, *35*, 1535.

[89] Chavanne, J. Radia (Version 4.29). [Online]. Available: `http://www.esrf.eu/Accelerators/Groups/InsertionDevices/Software/Radia` [2018, Sept 26], 2009.

[90] Elleaume, P.; Chubar, O.; Chavanne, J. *Proc. IEEE, PAC-97* **1997**, *3*, 3509.

[91] Chubar, O.; Elleaume, P.; Chavanne, J. *J. Synchrotron Radiat.* **1998**, *5*, 481.

[92] Verlet, L. *Phys. Rev.* **1967**, *159*, 98.

[93] Kramida, A.; Ralchenko, Y.; Reader, J.; and NIST ASD Team, NIST Atomic Spectra Database (Version 5.1). [Online]. Available: `http://physics.nist.gov/asd` [2018, Sept 26], 2013.

[94] Wiese, W. L.; Fuhr, J. R. *J. Phys. Chem. Ref. Data* **2007**, *36*, 1287.

[95] Radford, H. E.; Evenson, K. M. *Phys. Rev.* **1968**, *168*, 70.

[96] Beltrán-López, V.; Rangel G., J.; González-Nucamendi, A.; Jiménez-Mier, J.; Fuentes-Maya, A. *Phys. Rev. A* **1989**, *39*, 58.

[97] Holland, J. H. *Adaptation in Natural and Artificial Systems*; The University of Michigan Press, 1975.

[98] Igel, C.; Hansen, N.; Roth, S. *Evol. Comput.* **2007**, *15*, 1.

[99] Bates, D. M.; Watts, D. G. *Nonlinear Regression and Its Applications*; Wiley, 1988.

[100] Levenberg, K. *Quart. Appl. Math.* **1944**, *2*, 164.

[101] Marquardt, D. W. *SIAM J. Appl. Math.* **1963**, *11*, 431.

[102] Gilijamse, J. J.; Küpper, J.; Hoekstra, S.; Vanhaecke, N.; van de Meerakker, S. Y. T.; Meijer, G. *Phys. Rev. A* **2006**, *73*, 063410.

[103] Hayashi, K.; Sakudo, N.; Noda, T.; Takeda, A.; Fujimura, K.; Shimizu, N. *Nucl. Instr. and Meth. in Phys. Res. B* **1997**, *127-128*, 918.

[104] Scoles, G., Ed. *Atomic and Molecular Beam Methods*; Oxford University Press, 1988; Vol. 1&2.

[105] Kravchuk, T.; Reznikov, M.; Tichonov, P.; Avidor, N.; Meir, Y.; Bekkerman, A.; Alexandrowicz, G. *Science* **2011**, *331*, 319.

[106] Chang, Y.-P.; Długołęcki, K.; Küpper, J.; Rösch, D.; Wild, D.; Willitsch, S. *Science* **2013**, *342*, 98.

[107] Kirste, M.; Haak, H.; Meijer, G.; van de Meerakker, S. Y. T. *Rev. Sci. Instrum.* **2013**, *84*, 073113.

[108] Beardmore, J. P.; Palmer, A. J.; Kuiper, K. C.; Sang, R. T. *Rev. Sci. Instrum.* **2009**, *80*, 073105.

[109] Dulitz, K.; Softley, T. P. *Eur. Phys. J. D* **2016**, *70*, 19.

[110] Halbach, K. *Nucl. Instrum. Methods* **1980**, *169*, 1.

[111] Halbach, K. *Nucl. Instrum. Methods* **1981**, *187*, 109.

[112] Oldham, J. M. Combination of a cold ion and cold molecular source. Ph.D. thesis, University of Oxford, 2014.

[113] Rangwala, S. A.; Junglen, T.; Rieger, T.; Pinkse, P. W. H.; Rempe, G. *Phys. Rev. A* **2003**, *67*, 043406.