

Stark deceleration methods for cold molecule experiments

by

Yomay Shyur

B.A., Wellesley College, 2011

M.S., University of Colorado Boulder, 2014

A thesis submitted to the
Faculty of the Graduate School of the
University of Colorado in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
Department of Physics

2018

This thesis entitled:
Stark deceleration methods for cold molecule experiments
written by Yomay Shyur
has been approved for the Department of Physics

Prof. Heather Lewandowski

Prof. John Bohn

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Shyur, Yomay (Ph.D., Physics)

Stark deceleration methods for cold molecule experiments

Thesis directed by Prof. Heather Lewandowski

The growing field of cold molecule physics has demonstrated many methods of quantum control for precision measurement and molecular interaction studies. One particular method is Stark deceleration, which uses electric field gradients to produce packets of slow polar molecules. Since Stark deceleration does not increase phase-space density, the efficiency of the deceleration process is critical in determining the final cold molecule density. This dissertation describes three different Stark deceleration techniques. The first type of deceleration is pulsed-pin deceleration. We use this technique to produce slow OH molecules that can be trapped for cold collision studies. A two photon state-selective detection system, which includes a 118 nm photon, is used to determine the density distribution of OH in the trap. We explore the advantages and limitations to creating 118 nm light.

The second and third types of Stark deceleration are pulsed-ring and traveling-wave methods of running a ring-geometry decelerator. Ring-geometry decelerators present improved deceleration efficiencies due to the cylindrical symmetry of the electrodes and true three-dimensional confinement. The pulsed-ring mode of the ring decelerator uses commercial high-voltage switches to apply discrete pulses to the ring electrodes. The continuous operation mode of the ring decelerator uses varying sine-wave voltages to create a traveling Stark potential well and has proven to be incredibly challenging due to the analog high voltage requirements of the system. In order to bring ND₃ seeded in krypton (~ 415 m/s) down to rest in the laboratory, sine wave voltages with amplitudes up to ± 10 kV, currents up to 500 mA, and a bandwidth of 30 kHz down to DC are required. We describe a high-voltage amplifier that can be used for traveling-wave deceleration. The traveling-wave Stark deceleration mode is then compared to the pulsed-ring Stark deceleration mode and cases where one mode is superior to another are discussed.

Dedication

To my parents.

Acknowledgements

It is a pleasure to thank the people whose many efforts have made this dissertation possible. First, I would like to thank my advisor, Heather Lewandowski. She believed in me and entrusted me with an electronics project when I had very little background coming in. She always provided encouragement even when amplifier land seemed endless and bleak. She has been a patient mentor, who has taught me how to tell a story using data and brainstorm solutions to a myriad of scientific problems. All these skills will serve me well in future endeavors.

I also owe my gratitude to all the Lewandowski lab members I have worked with. Each one of you has taught me so much. Noah Fitch had the ambition to embark on the high-voltage amplifier journey, and provided valuable documentation and help even after moving to Europe. His groundwork in design, construction, and simulations for the ring decelerator have been critical to its operation. Travis Briles taught me a lot about optics and lasers and could provide references for seemingly endless cold molecule topics. Maya Fabrikant was always eager to help solve problems both in the lab and outside of lab. The construction of the ring decelerator and amplifiers would not have been possible without the work of David Macaluso, Kentaro Hoeger, and Ian Collett. I have also had the pleasure of working with John Gray and Jason Bossert, their eagerness to work on challenging experiments, persistence in finding the root of the problem, and ability to keep things upbeat in the lab have been invaluable. I wish both of them the best of luck as they continue their work on collision and reaction experiments. Additionally, help from and conversations with Philipp Schmid, James Greenberg, Kyle Miller, Cameron Straatsma, and Ben Saarel have been incredibly useful.

I have greatly benefited from discussions with the JILA cold molecule community. In particular, David Reens and Hao Wu have always provided keen insights into OH and Stark decelerator improvements. The construction, upkeep, repair, and upgrades to these mechanically and electronically complex decelerator experiments would not have been possible without the help of the JILA machine and electronics shops. Specifically, Terry Brown from the electronics shop deserves special recognition for the initial design of the amplifiers and many subsequent discussions.

Before embarking on my journey at JILA, I had many professors and mentors who encouraged me and guided me. I am thankful of the Wellesley College physics community where I have always found a home and Glenn Stark who gave me international physics research experience in my first year at Wellesley. Recognition also goes to Chris Arumainayagam, Juliet Pickering, Gillian Nave, and Gretchen Campbell who opened up their labs to me so that I could learn new things.

Thanks also goes out to Catherine Klauss, the best roommate ever, and my friends in Boulder for all the ski trips, cooking endeavors, encouragement, and support. And to my friends from afar, particularly Amanda and Rachel, who checked in regularly and were always interested in what was going on in the lab.

Lastly, I wish to thank my parents for their endless encouragement, without their support I would have never taken this journey.

Contents

Chapter

1	Introduction	1
1.1	Molecules in electric and magnetic fields	3
1.2	Beam Manipulation Techniques	5
1.3	Outline of this dissertation	7
2	Stark Deceleration	9
2.1	Applications of Stark Decelerators	9
2.1.1	Pin decelerators	10
2.1.2	Ring decelerators	12
2.1.3	Other Stark related techniques	14
2.2	Stark Decelerator Operation	15
2.2.1	Pulsed-pin Stark deceleration	15
2.2.2	Traveling-wave Stark Deceleration	24
2.2.3	Pulsed Ring Stark Deceleration	28
2.3	Phase-space Acceptance	28
2.4	Simulations	36
2.5	Experimental Layout	39
2.5.1	Supersonic Expansion	41
2.5.2	Detection	43

2.5.3	Our experiments	45
2.6	Decelerator comparison	52
3	Molecules of Interest	55
3.0.1	Rotational states	57
3.1	Ammonia	59
3.1.1	Stark effect in ND ₃	60
3.1.2	ND ₃ REMPI	61
3.2	OH	62
3.2.1	Hund's Coupling Cases	62
3.2.2	OH structure	64
3.2.3	OH in electric fields	65
4	Pulsed Ring Stark Deceleration	67
4.1	Pulsed-Ring versus traveling-wave Stark deceleration	67
4.1.1	Introduction	67
4.1.2	Experiment	71
4.1.3	Simple model of deceleration	74
4.1.4	Simulations	77
4.1.5	Experimental Results and Analysis	78
4.1.6	Conclusion	85
4.2	Implementing PRSD	86
4.3	Decelerating with PRSD	90
4.4	Transverse Dynamics	94
4.5	Calibration	97
4.6	8-ring periodicity PRSD	101

5	High-Voltage Amplifiers	105
5.1	High-Voltage Amplifier Electronics	106
5.1.1	Introduction	106
5.2	Design of the high-voltage amplifier	109
5.2.1	Leader Stage	111
5.2.2	Follower Stages	114
5.2.3	Feedback and control	116
5.2.4	Physical Construction	118
5.3	Amplifier Performance	121
5.3.1	Testing Setup	121
5.3.2	General Specifications	122
5.3.3	Voltage Sharing	127
5.3.4	Chirped sine waves	129
5.4	Conclusion	130
5.5	Individual amplifiers	130
5.6	Insulation	134
5.6.1	Amplifier bank	134
5.6.2	Amplifier walls	143
5.6.3	Optoisolator Failures	145
5.7	Operating multiple amplifiers	149
5.7.1	Digital and analog signals	149
5.7.2	Master Interlock	153
5.7.3	Enabling amplifiers	155
5.8	Amplifier Changes	158
5.8.1	Component value changes	158
5.8.2	Parts not in use anymore	162

6	Electrostatic trapping of OH	164
6.1	OH experimental setup	167
6.2	OH REMPI Detection	172
6.3	OH 1+1' REMPI with 118 nm photons	173
6.4	OH in electric and magnetic fields	177
6.4.1	Hamiltonian for OH in electric and magnetic fields	179
6.4.2	Electrostatically trapped OH with a uniform magnetic field	180
6.4.3	Future work	185
6.5	118 nm light	186
6.5.1	Creating 118 nm light	187
6.5.2	Experimental setup	190
6.5.3	Limitations to 118 nm light production	191
6.5.4	Possible explanations for the limitations	197
	Bibliography	203
	Appendix	
A	Circuit Diagrams	218
B	Modifications summary	245
B.1	Amplifier bank	245
B.2	Amplifier breakout boards	246
B.3	Control boards	246
B.4	Master interlock	246
B.5	Insulation	246

Tables

Table

2.1	Typical values for the full width at half maximum (FWHM) of the initial distribution of molecules	37
2.2	Properties of the two decelerators used.	54
3.1	Properties of ND ₃	59
4.1	Switching rise times for PPSD vs PRSD	90
5.1	Technical parameters of the amplifier.	123
5.2	Small-signal amplifier parameters. Rise and fall times, t_r and t_f , and delay d were taken with a 1 kHz square wave with an output of V on a load capacitance of C_l . The rails were held fixed at ± 11 kV.	126
5.3	RTVS 3-95-1 properties	139
5.4	Changes to amplifier boards resistor values	159
5.5	Changes to amplifier breakout boards	159
5.6	Changes to control boards	160

Figures

Figure

2.1	End on view of the Stark decelerators in our laboratory	15
2.2	Molecule in a dipole field	16
2.3	Potentials created in a pin decelerator	17
2.4	Timing molecule potentials for $\phi_0 = 0^\circ$	18
2.5	Phase angles on a longitudinal potential	19
2.6	Potentials used for various phase angles	21
2.7	Position of molecules ahead and behind the timing molecule when the potentials switch	23
2.8	CAD drawings of one ring electrode and the full decelerator	25
2.9	TWSD voltages on ring electrodes and the resulting potentials	27
2.10	Energy lost, ΔK per stage	31
2.11	The longitudinal (z axis) separatrices for $\phi_0 = 0^\circ, 20^\circ, 40^\circ, 60^\circ, 80^\circ$	32
2.12	The effective longitudinal potential for an initial velocity of 415 m/s and various final velocities	34
2.13	Normalized longitudinal potentials for TWSD, PRSD, PPSD, and a cosine	36
2.14	PRSD acceleration as a function of phase angle	36
2.15	Experimental setup for a general Stark deceleration molecule	40
2.16	Simple model of LIF and REMPI detection	44
2.17	Time-of-flight measurements for both the PZT and NVP valves	46
2.18	A CAD drawing of the pulse-pin decelerator	48

2.19	Electrostatic trap used with the pulsed-pin Stark decelerator	48
2.20	The longitudinal on axis potentials for the trap on (blue) and half on (red) configurations of the trap	49
2.21	Physical assembly of a ring decelerator	51
2.22	The ring trap	52
2.23	On-axis longitudinal potentials for PPSD, PRSD, and TWSD modes in our pin and ring decelerators	53
3.1	The rotational levels of ND_3 in the electronic and vibrational ground state	60
3.2	The Stark energy as a function of electric field for ND_3 and NH_3	61
3.3	The 2+1 REMPI scheme for detecting ND_3	62
3.4	The rotational structure in the electronic and vibrational ground state of OH	65
3.5	The Stark splitting in OH	66
4.1	A schematic of the experimental setup for PRSD	71
4.2	Ring number, applied voltages, positions of electrodes, and longitudinal Stark potentials of ND_3 for TWSD and PRSD	73
4.3	The instantaneous TWSD and PRSD transverse Stark potentials in the x-direction for the voltage configuration in Fig. 4.2 and at $z=0$ mm	75
4.4	The PRSD and TWSD Stark potential wells with applied linear fictitious potentials for various accelerations	76
4.5	Time-of-flight traces of a molecular beam decelerated with PRSD	79
4.6	Time-of-flight traces of a molecular beam decelerated with TWSD	80
4.7	ND_3 signal verses acceleration for both TWSD and PRSD	81
4.8	Longitudinal phase-space distributions for TWSD and PRSD for an initial velocity of 415 m/s and final velocities of 320 m/s and 200 m/s	82
4.9	Simulated integrated ND_3 signal verses acceleration for both TWSD and PRSD	83
4.10	The acceleration at the crossing point of simulated TWSD and PRSD signal	85

4.11	Electrode voltages for PRSD	86
4.12	Recommended Behlke control circuit	89
4.13	PRSD at ± 7.5 kV	91
4.14	LPSA for PRSD at different peak voltages	91
4.15	LPSA for PRSD at ± 7 kV	93
4.16	LPSA for PRSD at ± 10 kV	94
4.17	Stark Potential energy surface $xz(y=0)$ plane	95
4.18	Transverse potential for TWSD	95
4.19	PRSD and TWSD transverse potential progression	97
4.20	Time-of-flight uncalibrated and calibrated PRSD scans	100
4.21	Electrode voltages for 8R PRSD	103
4.22	Comparison of 4R PRSD and 8R PRSD simulations	104
5.1	An optically coupled, all N-channel push-pull amplifier	110
5.2	A leader stage of the high-voltage amplifier	112
5.3	A follower stage of the high-voltage amplifier	115
5.4	A photograph of the amplifier with the top lids removed	120
5.5	The test setup for measuring the output voltage and current	122
5.6	Measured gain of the amplifier	123
5.7	Measured output sine waves and their respective Fourier transforms for 4 and 30 kHz input sine waves	125
5.8	Ratio of voltage dropped between the final follower stage (F9) and the leader stage (L) on the bottom bank for three different frequencies	127
5.9	A linearly chirped voltage waveform from 30 kHz to DC at ± 10 kV	129
5.10	The high-voltage amplifier with its probe and breakout box	132
5.11	Amplifier control box front panel	132
5.12	Setup for individual amplifier testing on the bench	133

5.13	Probing an amplifier bank	133
5.14	± 8 kV output without and with the amplifier box filled with SF ₆	135
5.15	Persistent current oscillation	136
5.16	± 12 kV output voltage at 5 kHz	138
5.17	A thin layer of RTVS 3-95-1 potting material on one amplifier bank	140
5.18	Fully potted amplifier banks	140
5.19	Cross section of custom high voltage cable	143
5.20	Current dips caused by proximity to polycarbonate	143
5.21	Amplifier box polycarbonate walls	144
5.22	Pictures from an amplifier that had a short	146
5.23	Shorting out the current monitoring optoisolator	147
5.24	A simplified schematic of amplifier electronic and control connections	148
5.25	Digital timing sequence	153
5.26	Decelerator rod voltage with grounded and non-grounded amplifier input	158
5.27	Control board rail monitor circuit	161
6.1	The gas mixing system, valve, and discharge plates for creating a supersonic beam of OH	167
6.2	The detection region configuration for detecting trapped OH	168
6.3	The length in time versus the stage number for one mixed deceleration timing sequence	169
6.4	A diagram of two different OH trap loading schemes	171
6.5	An energy level diagram of OH showing the relevant states for 1+1' detection	173
6.6	The central density OH in the trap for two different loading schemes	175
6.7	The FWHM of trapped OH over time	175
6.8	The width of the OH cloud as it expands after the trap is turned off	177
6.9	The electric field in the trap along a longitudinal (a) and transverse line (b)	178

6.10	The electrostatic trap with Helmholtz coils to produce a uniform magnetic field at the center of the trap	178
6.11	OH energy levels in a uniform 90 G magnetic field with varying β and electric field	181
6.12	OH energy levels in a uniform 500 G magnetic field with varying β and electric field	182
6.13	OH energy levels in a uniform 800 G magnetic field with varying β and electric field	183
6.14	The region in the x-z plane ($y=0$) of our trap where the electric field is greater than 6 kV/cm is shown in red	184
6.15	Plot of $ F_1(b\Delta k) ^2$ in the tight-focusing limit and for $b = 1$	188
6.16	A diagram of the 118 nm gas mixing cell and the ionization and detection region	191
6.17	Phase-matching curves at different pressures of Xe	193
6.18	A picture of the fluorescence streak that appears near the 355 nm focal point	193
6.19	A phase-matching curve, with $\mathcal{P}_{Xe} = 50$ torr, taken while monitoring both the 118 nm flux and the fluorescence intensity	194
6.20	A comparison of the 118 nm signal and fluorescence streak signal at different xenon pressures for a phase-matched mixture with Ar (Ar:Xe Ratio = 11.3) and Kr (Kr:Xe Ratio = 4.6)	195
6.21	A spectrum of the fluorescence streak created in a xenon-argon mixture	196
6.22	A spectrum of the fluorescence streak created in a xenon-krypton mixture	197
6.23	A cartoon of the 118 nm gas mixing cell, absorption cell, and ionization and detection region	199
6.24	Diagram depicting the feedback mechanism of electric field induced SFG	201

Chapter 1

Introduction

Pushing the limits of quantum control of atoms has led to the discovery of new states of matter [1, 2] and a greater understanding of quantum phenomena. Studies of cold atoms have led to new discoveries in fundamental physics [3], precision measurement [4–6], and quantum processes [7, 8]. The field of cold molecule research is a natural and challenging extension of cold atom research. Currently, several experimental techniques have been able to cool down molecules to the millikelvin temperature range, and select techniques using specific molecules have reached the microkelvin range and lower. Molecules have many more degrees of freedom due to their ability to vibrate and rotate, and these motions can couple together to provide a rich spectroscopic landscape. However, these additional degrees of freedom make molecules more complex and more difficult to cool than atoms.

Cold molecules can be used for spectroscopic, precision measurement, collision, and reaction studies. Low temperature spectroscopy provides improved spectral line resolution and limits spectral congestion, since fewer quantum states are populated [9]. Cold molecules can be used as precision measurement tools and tests of fundamental physics, including studies of the fine structure constant and proton-to-electron mass ratio [10, 11] and the electron electric dipole moment [12, 13]. At cold temperatures, applied electric or magnetic fields can be used to control the energies of the reactants in order to form a better understanding of chemical reactions. These fields can be helpful for control of molecular orientation [14, 15] and stereochemistry studies [16, 17]. Cold molecules provide the opportunity to gain control over molecule collisions and reactions. At low tempera-

tures, interactions are dominated by quantum mechanical effects. Tunneling and resonances are important to the understanding of chemical dynamics when the energy available is lower than the activation energy of the process. By “freezing” out certain energy levels and restricting molecules to one or a few quantum states, state selective processes can be studied. One goal of these studies is full quantum control over a reaction, including well-defined states of the reactants and products, which will allow for a true quantum understanding of basic chemical reactions.

The field of cold molecules encompasses all techniques used to create cold molecules, many of which are tailored to specific categories of molecules. Here, we provide an overview of many of the techniques that have been demonstrated and explored in just the last twenty or so years in order to provide a context for the main technique discussed in this dissertation, Stark deceleration. However, this is by no means an exhaustive list and the reader is encouraged to consult the plethora of review articles have been written on cold molecules [18–27].

There are two methods of creating cold molecules, indirect and direct. Indirect methods take cold atoms and associate them to form molecules. This can be accomplished by photo-association [28, 29] or magneto-association [30] and forms molecules that are ultracold, in the hundreds of nanokelvin range [31]. These techniques generally can produce molecules that are much colder and with higher densities than direct techniques, but they are limited to alkali and alkali-earth atoms that can be laser cooled before association. Direct methods, on the other hand, take molecules from near room temperature and make them cold. They are applicable to many atmospheric and astrophysically relevant species such as NH_3 , OH , NH , and CH . For the most part, the molecules can only be made cold, in the range of a few degrees Kelvin to millikelvin, and not ultracold. However, recent developments in laser cooling of molecules have reached temperatures in the microkelvin range [32, 33]. Direct methods are the focus of the remainder of this chapter.

Within direct cooling methods, there are phase-space compressing methods and non-phase space compressing methods. Phase-space compression represents true cooling, where an irreversible process has removed entropy from the system. These techniques include RF evaporation [34], optoelectrical Sisyphus cooling [35], buffer gas cooling [36–38], and laser cooling of molecules [32,

33,39–42]. While laser cooling and magneto-optical trapping have been fundamental for cold atom studies, they rely on cycling transitions that allow atoms to scatter light on the same transition many times. Molecules have complex energy structures due to their additional degrees of freedom, which severely restricts the options for cycling transitions and thus laser cooling molecules. Nonetheless, laser cooling of diatomic molecules has been demonstrated with SrF [39], YO [40], CaF [41], and YbF [33], and these molecules can also be loaded into MOTs [32,42].

Non-phase-space-compressing methods involve filtering or slowing and generally result in low densities of molecules. These non-compressing methods typically begin with a room temperature molecular beam, but could also be loaded from a cryogenic source [43]. Many of these molecules of interest are highly reactive and must be created *in-situ* using either photodissociation or an electric discharge. While precision control of molecular beams is an established technique [44–46], directly cooling these beams is a more recent development [24]. These experiments are generally not in thermal equilibrium, and the description of “cold” is a measurement of the kinetic energy stored in the degrees of freedom and not a measurement of temperature. For example, a supersonic expansion produces a “cold” beam. It has molecules that are internally cold (close to the rovibronic ground state) and with a narrow velocity distribution. However, this beam may still be moving fast (300-600 m/s) in the laboratory frame. Non-phase-space-compressing direct techniques can then be used to slow down the beam and reduce the average velocity. These techniques are able to slow down the molecules in the laboratory frame, but they do not make the sample “colder.”

1.1 Molecules in electric and magnetic fields

Externally applied electric and magnetic fields are commonly used to control atoms and molecules. While it is common to think of controlling ions in this fashion, it may not be immediately clear that this is a useful technique to use on neutral molecules. However, if a molecule has an electric (magnetic) dipole moment, it will experience an energy shift in the presence of an electric (magnetic) field due to the Stark (Zeeman) effect. If the field is inhomogenous the changing energy shift will be experienced as a force. This force serves as a handle to manipulate the molecules and

allows them to be accelerated, decelerated, and trapped.

Direct cooling methods, which decelerate molecules from a molecular beam, generally exploit inhomogeneous fields as the decelerating force. To motivate the need for inhomogeneous fields and the importance of field gradients, we examine the behavior of an electric dipole in an external field. In the classical picture, the Hamiltonian for the interaction of a dipole moment with a electric field is

$$H_{field} = -\mu_e \cdot \mathbf{E}, \quad (1.1)$$

where μ_e is the electric dipole moment and \mathbf{E} is the electric field. This gives the potential energy of the system to be

$$U = -\mu_e \cdot \mathbf{E}, \quad (1.2)$$

and the force to be

$$\mathbf{F} = -\nabla U = \vec{\mu}_e \cdot \nabla \mathbf{E}. \quad (1.3)$$

Equation 1.3 indicates that molecules with electric dipole moments can be controlled using inhomogeneous fields since the gradient of the fields sets the force applied to the molecules. As will be seen extensively in Chapters 2 and 4, the gradient of the field and how it is applied are the most important aspects to the success of any deceleration method.

In this classical picture, the dipole will minimize the potential energy by aligning with the field. However, in the quantum mechanical picture, the dipole moment cannot freely rotate and its projection on to the applied field is quantized. For the case of the potential energy (Equation 1.2)

$$U = -\mu_e \cdot \mathbf{E} = -|\mu_e||\mathbf{E}| \cos(\theta), \quad (1.4)$$

where θ is the angle between μ_e and \mathbf{E} , and both $|\mu_e|$ and θ are fixed, this leads to the sign of U being dependent on θ . The molecule's energy either increases or decreases when a field is applied. Molecules whose energies increase with the field are referred to as weak-field-seeking, and those whose energies decrease with the field are referred to as strong-field-seeking. Strong-field-seeking states are attracted to regions of strong fields, which are found near the surfaces of electrodes. This

makes them difficult to work with since it is not possible to confine strong-field-seeking molecules in any static potential. Instead, it is much easier to work with weak-field-seeking molecules since it is possible to create areas of low electric fields in free space.

These same arguments can be made for a magnetic dipole in a magnetic field by replacing μ_e with the magnetic dipole moment μ_b and the electric field \mathbf{E} with the magnetic field \mathbf{B} . While there are many similarities between the interactions of a dipole and an electric and magnetic field, the key difference is that the magnetic dipole is permanent and always exists in a molecule, and the electric dipole moment is induced and exists only when a field is present.

1.2 Beam Manipulation Techniques

While this dissertation uses Stark deceleration of weak-field-seeking states to create cold molecules, many other beam manipulation techniques have been used to decelerate molecules. A detailed description of deceleration of weak-field-seeking states and the history of the technique is provided in Chapter 2.

The complement to standard weak-field-seeking Stark deceleration is alternate gradient deceleration for strong-field-seeking molecular states [47–50]. This technique has been used to decelerate CO [47], OH [51], YbF [48], and benzonitrile ($\text{C}_7\text{H}_5\text{N}$) [52]. However, alternate gradient deceleration is not as stable as weak-field-seeking methods since it is highly dependent on the alignment of the decelerator and longer decelerators have rapidly decreasing efficiencies. Also, using the Stark effect, Rydberg-Stark deceleration can be used to decelerate long-lived Rydberg states [53, 54], which are readily polarizable even if the ground state molecule is not. Rydberg-Stark deceleration has also been demonstrated on chips [55] and decelerated H-Rydberg state molecules have been brought to a stop [56].

Apart from electric fields, similar inhomogeneous field techniques have been demonstrated with magnetic fields. Zeeman deceleration has been used with atomic hydrogen [57, 58] and neon [59]. Zeeman deceleration offers the ability to decelerate non-polar free radicals, which cannot be done in a Stark decelerator.

Other methods of decelerating molecules from a beam, which do not make use of inhomogeneous fields, have also been constructed. These include filtering, collision based schemes, optical methods, and physical translation methods. Quadrupole velocity filtering uses a bent quadrupole to select slow molecules out of an effusive beam [60–62]. Crossed-beam kinematic cooling brings molecules to rest using one collision with a particle of carefully chosen mass and velocity such that the collision partner’s momentum vector is opposite that of the molecule to be cooled [63,64]. Optical methods include carefully selected photodissociation, such that the recoil speed of the fragment of interest equals the molecular beam speed [65]. Additionally, optical Stark deceleration uses time-varying red-detuned optical lattices [66], and has been demonstrated experimentally using nitric oxide [67] and benzene [68]. Both pulsed [69] and traveling-wave versions [70] of optical Stark deceleration have been proposed.

The last approach involves physical rotation of the molecular source or decelerating apparatus in order to change the velocity of the molecules in the laboratory frame. One technique places the molecular beam source on a rotating stage such that it moves backwards during the expansion [71,72]. The beam out of the translating source is slower in the lab frame than had the source been stationary. Another technique uses a novel centrifuge decelerator [73]. It uses a spinning quadrupole to guide and decelerate molecules up a centrifugal potential energy hill toward the central axis. This technique has the benefit that it can produce truly continuous sources of slow molecules if the initial source is continuous. The centrifugal decelerator has also been loaded from a cryogenic buffer gas cell in a system called a cryofuge, which has been used to decelerate fluoromethane and deuterated ammonia [74].

Once the molecules are decelerated, the slow beam can be used directly in crossed-beam experiments [75–78], used with other further deceleration techniques [79], loaded into molecular storage rings [80,81] and synchrotrons [82,83], or loaded into a trap. Molecule traps can include electrostatic [84–86], magnetic [87,88], and ac traps [89–91], and they can be used in beam-trap [92,93] and cotrapped experiments [94].

The last category of experiment, cotrapped experiments, allows for longer interaction times

of the trapped species. Although trap densities are generally not large enough to see reactions, collisions with a separate denser sample can yield important measurements of collisional cross sections [94]. Within a cotrapped experiment, there are two types of collisions, elastic and inelastic. In elastic collisions, the internal energies of the constituents do not change and only momentum is exchanged between the collision partners. On the other hand, in inelastic collisions, internal quantum states of the particles change. Since only specific quantum states can be held in a trap, inelastic collisions will lead to trap loss and a decreased density of the trapped sample. If elastic collisions occur, then two samples of particles with different temperatures can thermalize in the trap. This allows for sympathetic cooling, one goal of cold molecule research that has yet to be demonstrated.

1.3 Outline of this dissertation

This dissertation addresses several improvements to Stark decelerated cold molecule experiments, including improvements to molecule deceleration and molecule detection. These techniques come together in the last chapter in a discussion of cotrapped hydroxyl (OH) and rubidium. Three different Stark deceleration techniques are utilized in this dissertation. The first technique is a discrete mode, where set voltages are pulsed on and off on a decelerator made up of crossed pin electrodes. This mode is referred to as pulse-pin Stark deceleration (PPSD). The second and third techniques are both performed on a ring electrode decelerator. The second technique is a discrete mode and is referred to as pulsed-ring Stark deceleration (PRSD). The third technique is a continuous mode, called traveling-wave Stark deceleration (TWSD), where analog smoothly-varying voltages are applied to the ring electrodes. Chapter 2 discusses all three types of Stark deceleration and gives an overview of the decelerators in our lab. Chapter 3 covers the main molecules, deuterated ammonia (ND_3) and OH, that are used to test these deceleration techniques. Both methods of running a ring-geometry decelerator are discussed in Chapter 4. This includes a comparison between the traveling-wave and pulsed-ring methods, how to determine when one method is superior, and how to implement the pulsed mode of the ring Stark deceleration. In order to implement the

traveling-wave mode of a ring decelerator, analog high-voltage amplifiers are required. Amplifiers that can produce the required voltage waveforms for the ring decelerator in our lab were first presented in Noah Fitch's dissertation [95]. In Chapter 5, we discuss the required modifications for both components and the construction of the amplifiers in order to keep them operating successfully. A complete set of updated circuit diagrams for the amplifiers can be found in Appendix A and a full list of changes between the circuits presented in Noah's dissertation and this dissertation can be found in Appendix B. Lastly, an experiment that blends Stark deceleration and cold trapped atoms is discussed from molecule creation to detection. Cold electrostatically trapped OH and ultracold magnetically trapped rubidium atoms provide an interesting landscape for studying the effects electric and magnetic fields on OH and collision cross sections for the system. Here, a modified method of running a pulsed-pin decelerator and an improved method of state-selective detection of OH molecules are discussed.

Chapter 2

Stark Deceleration

Stark deceleration uses spatially inhomogeneous and time-varying electric fields to decelerate neutral polar molecules resulting in a beam of slow molecules with a tunable average final velocity. The molecules experience a force, due to the Stark effect, which will increase or decrease the energy of a quantum state as a function of the strength of the electric field. Quantum states where the energy increases with the strength of the electric field are known as weak-field-seeking states, and those where the energy decreases with the strength of the electric field are known as strong-field-seeking states. The changing Stark energies due to inhomogeneous electric fields are exploited to create a slow beam of molecules by arranging electric fields such that a beam of molecules must always climb up a potential energy hill. In standard Stark deceleration, only the weak-field-seeking states can be decelerated in this manner. This chapter provides an introduction to the development and uses of Stark decelerated beams and is followed by a detailed description of three different types of Stark deceleration. We then discuss the phase-space acceptance calculations, molecular trajectory simulation of molecules in the decelerator, and, lastly, experimental setup of the decelerator experiments. The subsequent chapters will cover many of these topics in more depth.

2.1 Applications of Stark Decelerators

Stark decelerated beams are advantageous for collision and crossed-beam experiments since the final velocity of the beam, and thus the energy can be carefully controlled [75]. Decelerated

beams, usually less than 100 m/s, have been used in used in crossed beam [77, 78, 96, 97], molecular fountain [98], beam-trap [92, 93], and cotrapped [94] experiments. These various molecule interaction configurations can be used to study high-resolution spectroscopy [10, 98, 99], collision physics [100], radiative lifetimes of molecules [101, 102], blackbody optical pumping rates out of trapped states [103], and tests of fundamental physics [10, 104, 105].

Stark deceleration is a useful tool for creating velocity tuned, cold and slow beams of molecules. Molecular packets are introduced to the Stark decelerator with pulsed valves. The packet is internally cooled during the supersonic expansion, but maintains a high forward velocity, the deceleration process then slows down the beam. There is no phase-space compression during deceleration since it is a conservative process. Stark deceleration does not produce a pure quantum state beam [106], instead other filtering techniques are required if a completely pure beam is desired. The only ways to increase the signal, the number of slow and detected molecules, are a better source, more efficient deceleration, or more efficient detection (improved signal-to-noise). Improvements to the source and detection are discussed in Section 2.5 and below we briefly cover the developments and improvements to Stark deceleration since its first experimental realization.

2.1.1 Pin decelerators

The first Stark decelerator, which we refer to as the pulsed-pin Stark decelerator, was constructed in 1999 by Bethlem *et al.* [107]. It had 63 stages and was used to decelerate CO from 225 m/s down to 98 m/s. Since then, the molecules NH₃/ND₃ [84], OH [108], H₂CO [109], NH [76], LiF [110], SO₂ [111, 112], CaF [50], and NO [113] have been decelerated using PPSD.

The molecules from a pulsed-pin Stark decelerator can be loaded into electrostatic and magnetic traps. The type of trap will depend on the molecule, type of study, and how much optical or physical access to the trap center is required. For example, the simplest electric quadrupole trap consists of three electrodes [84], but does not have easy access for a detection laser. To remedy this, a small hole can be drilled in the center electrode for optical access [84] or the center electrode can be split into two electrodes [85]. Both these solutions slightly alter the trap from a

perfect quadrupole. Thus far NH_3/ND_3 [84, 85, 114], NH [115], and OH/OD [86, 103] have been trapped electrostatically and OH [87] and NH [88] have been trapped in magnetic traps. A high-field-seeking state of $^{15}\text{ND}_3$ has been trapped with an ac trap after the low-field-seeking state was decelerated and the molecules were pumped via a microwave transition into the high-field-seeking state [49, 89, 91].

The majority of Stark deceleration experiments have been conducted with beams from PPSD. Beam and trapped PPSD experiments are often limited by the low numbers of successfully decelerated and detected molecules. While careful optimization of beam production and decelerator loading can improve the density of the decelerated beam [116], the density at the end of the decelerator is still limited by the initial gas pulse, decelerator dynamics, and molecule detection efficiency. Additionally, there are distinct differences between the phase-space output of the pin decelerator and the phase-space acceptance of any trap placed after the decelerator which will further reduce the efficiency.

Many efforts have been made to improve efficiency of a pin decelerator [117–120]. Various timing schemes and operating modes can improve specific aspects of the performance of the pulsed-pin decelerator, however the crossed-pin electrode geometry fundamentally limits the efficiency due to longitudinal and transverse coupling. Alternative deceleration methods include not using all the stages for deceleration [117, 119] and adjusting how many times the decelerator switches in comparison with the number of deceleration stages [118, 120]. In the first category, a subset of the decelerator stages are used for deceleration and the remaining stages are used for bunching, where the velocity of the beam does not change but the molecules are “bunched” together in phase space. By using more aggressive deceleration in a subset of stages, the resulting packet has a narrower velocity spread due to the smaller phase-space acceptance at higher deceleration phase-angles. This can be accomplished by either bunching at the beginning [119] or bunching at the end [117]. While a similar packet could be obtained with a shorter decelerator, modifying the length of a decelerator is time consuming. Additionally, if the decelerator is to be used with multiple types of molecules or carrier gases, building the decelerator to the longest required length and then using alternative

deceleration sequences, when changing to different molecules, will save the most time. Bunching at the beginning yields slightly more molecules in the decelerated packet for $v_f < 50$ m/s because there are fewer losses when bunching at higher velocities compared to lower velocities since less time is spent in the decelerator. Bunching at the end of the deceleration increases the separation in time between the undecelerated packet and the decelerated packets, which can be useful for collision experiments.

In the second category of alternate deceleration methods, the number of timing switches is altered and not equal to the number of pin stages [118, 120]. One mode of operation switches the voltage configuration once for every three sets of pins the timing molecule traverses. In this configuration, only every third set of pins is used for deceleration, the other two sets of pins provide transverse focusing. This method is known as S=3, where S indicates the stage periodicity for deceleration. Traditional PPSD, where every stage decelerates, is referred to as S=1 deceleration. The reduction in stages used for deceleration means that either a longer decelerator is required or less deceleration is achieved. Scharfenberg *et al.* [118] found that for a 2.6 m long decelerator and a beam of OH molecules with an initial velocity of 350 m/s, S=3 deceleration is more efficient than S=1 down to a final velocity of 100 m/s. Although S=3 is more efficient at moderate final velocities, at lower final velocities, S=1 is more efficient, since fewer pin pairs result in molecules at low velocities spend less time in the decelerator. The S=3 and S=1 modes have also been combined such that S=3 is used initially and then the decelerator is switched to S=1. This must be completed before the molecular packet reaches the cross over threshold velocity around 100 m/s. Additionally, methods have been used where S=3 and S=1 methods are intermixed (see Section 6.1).

2.1.2 Ring decelerators

The first use of macroscopic traveling-wave Stark potentials with a ring-geometry decelerator was in 2010 by Osterwalder *et al.* [121] using a beam of CO, which was decelerated from 288 m/s to 144 m/s. While TWSD has many advantages, it is more challenging to implement than PPSD because it requires high-voltage analog amplifiers with demanding specifications to generate the

moving Stark potential wells. Early TWSD experiments were limited by the lack of availability of high-voltage amplifiers with a large bandwidth. The amplifiers used to generate the voltages for the electrodes in the first TWSD experiment operated at ± 8 -10 kV with 10-30 kHz sinusoidal voltages [121, 122]. The limited bandwidth of the amplifiers meant that the decelerator could achieve only moderate deceleration and the molecules could not be brought to rest.

Commercial amplifiers (*Trek 5/80*), which are capable of a larger frequency range (0-60 kHz) and outputs up to ± 5 kV, make it impossible to decelerate a supersonic beam down to rest with a decelerator that is only one meter long. For example, a beam of SrF seeded in Xe with an initial velocity of 290 m/s would require a ~ 5 m long decelerator [123]. However, amplifiers, such as the *Trek 5/80*, with limited maximum voltages are particularly useful for the TWSD of heavy molecules used in precision measurements, since the electric fields produced are low enough to keep the molecules in the weak-field seeking regime, where the Stark energy increases with increasing electric field [124]. To date, YbF [125] and SrF [123] have been decelerated using TWSD and the longest demonstrated decelerator (4 m) slowed SrF from 290 to 120 m/s using ± 5 kV amplifiers [126].

A combination deceleration method has also been used, to take molecules to rest and trap them within the ring electrodes. It couples a traveling-wave decelerator to the end of a pulsed-pin decelerator [127]. In this combination technique, the molecules are decelerated in the pin decelerator to velocities around 100 m/s, and then enter the traveling-wave Stark decelerator where they are brought to rest. The molecule packet can also be temporarily held inside the ring decelerator. This coupled deceleration method combines the simplicity of PPSD with the low velocity advantages of TWSD. It takes advantage of the cylindrical symmetry of the ring geometry at low velocities, but it relaxes the requirement to have either extremely high maximum voltages or a very long decelerator to reach electrostatically trappable velocities. Using this combination technique, NH_3 , ND_3 , and CH_3F have been trapped, manipulated, and adiabatically cooled within the ring decelerator [127–130].

Deceleration of a supersonic beam moving at 300-450 m/s to rest using only TWSD in a ring

decelerator has been challenging to achieve due to electronic and hardware requirements and has not been completed to date. However, many strides have been made toward this goal and high voltage amplifiers to help achieve this are described in Chapter 5.

2.1.3 Other Stark related techniques

To date, pin-geometry decelerators have been the dominant form of Stark deceleration implemented in the lab. Only a few groups work with ring-decelerators. In addition to rod and ring decelerators, a few other decelerator geometries have been proposed. One has sets of ‘U’ shaped electrodes and transverse charged confining wires [131]. While similar to the pin configuration, this particular geometry has true potential wells since each stage has transverse confinement in two axes. In this configuration the voltages required for transverse confinement are very high (26 kV), and thus is very challenging to build.

Several Stark deceleration techniques have also been demonstrated in miniature with microelectrodes on a chip. Meek *et al.* [132] first demonstrated the idea for spatially modulated potential wells on a microstructured chip used to guide and then decelerate and trap metastable CO molecules [133, 134]. These systems have been used to study how magnetic fields can be used to suppress non-adiabatic losses of electrostatically trapped CO [135] and to manipulate the temperature of trapped CO [136].

While only weak-field-seeking molecules have been discussed so far, deceleration techniques for strong-field-seeking molecules have also been developed with limited success. These alternate gradient decelerators have been used to decelerate CO [90, 114] and YbF [48].

Finally, a few other useful beam manipulation tools can be used in addition to decelerators. Hexapoles and additional crossed rods can be used for longitudinal focusing and cooling [137] and molecule storage rings/synchrotrons of various geometries [80–83] can be placed after the decelerator. A recent demonstration of a beam merger made of electroplated 3D printed plastic [138] opens up exciting new possibilities in a variety of macroscopic beam guide and manipulation techniques. This technique allows for easy construction of electrode geometries that were previously

too difficult or impossible to machine. Selective electrode plating of a 3D printed object allows for uniquely shaped electrodes and electrode supports in ways that were previously not possible.

2.2 Stark Decelerator Operation

Decelerators come in two main geometries, pins and rings, Figure 2.1. The pin decelerator is always run with high-voltage pulses, while the ring decelerator can be run with high-voltage pulses or high-voltage sine waves. The decelerators will be discussed in the order that they were experimentally realized.




Figure 2.1: End on view of the two Stark decelerators in our laboratory.

2.2.1 Pulsed-pin Stark deceleration

A molecule in a weak-field-seeking state near a strong electric field created by two charged rods sees a potential energy hill, Figure 2.2. As the molecule travels into the electric field, it exchanges kinetic energy for potential energy. If the molecule is allowed to travel past the peak of the potential energy hill, it will accelerate back to its initial velocity. However, if instead, the electrodes turn off right as the molecule reaches the peak of the hill, kinetic energy in the forward direction will have been removed from the molecule. This sequence can be repeated over and over in order to removed significant forward velocity and decelerate a molecule.

The potential energy hills are created by sets of crossed pins, typically mounted in four sets

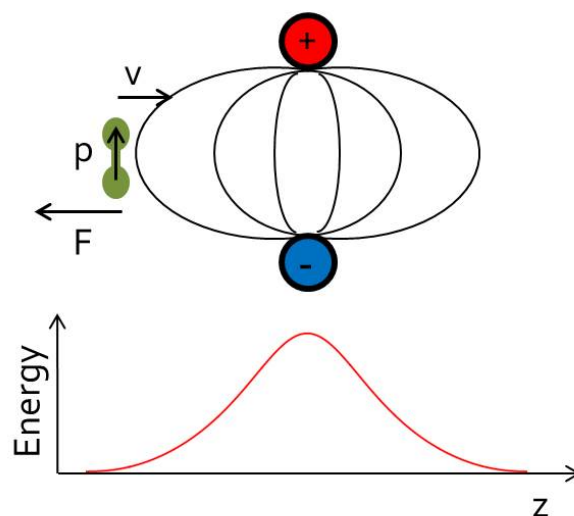


Figure 2.2: A weak-field-seeking molecule (green) with a dipole moment p and velocity v approaches a non-uniform electric field. Due to the Stark effect, the molecule sees the high-field region as a potential energy hill and experiences a force opposite its velocity.

of rods running the length of the decelerator (Figure 2.1). Pins are then placed into each rod at set distances such that when the four rods are assembled there are alternating sets of crossed pins. In the physical decelerator, the electrode pairs are oriented perpendicular to each other. This combination of rods and pins simplifies the electronics required to run the Stark decelerator. All pins on one rod have their voltage switched simultaneously, and only four high voltage switches are required to run the entire decelerator.

Figure 2.3 shows how applying a voltage to alternating sets of electrodes produces periodic Stark potential wells that shift longitudinally. For this series of pin electrodes (depicted as circles), only two voltage configurations are required to create a time-averaged, over all the deceleration stages, moving potential well. The longitudinal potential for each voltage configuration is given immediately below the pin diagram. The two configurations are almost identical, but the potential well shifts longitudinally by half a period between the two configurations. By alternating between these two configurations at the correct times, a molecule can be decelerated.

To determine the correct timing sequence for the voltage configurations, we will initially consider only one molecule traveling along these longitudinal Stark potential wells and always

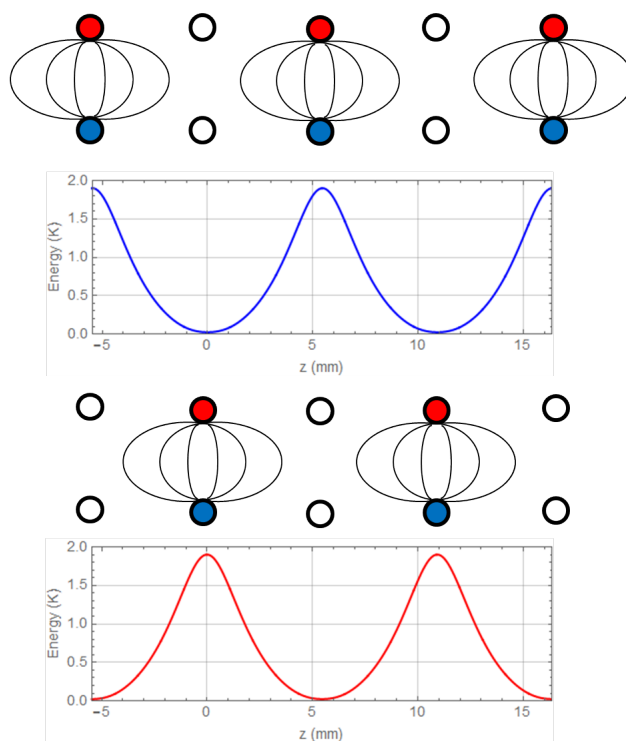


Figure 2.3: The two voltage configurations in a pin decelerator and the potentials they create. In each voltage configuration, red corresponds to positive high voltage, blue to negative high voltage, and white to ground. The Stark potentials are plotted below each voltage configuration. The values for the potentials are for ND_3 in the pin decelerator in our lab operated at ± 12 kV.

confined to the central axis of the decelerator. In this dissertation, the decelerator axis will be referred to as the z -axis. This molecule that travels only on the z -axis with no transverse (x or y) motion is called the timing molecule. While in practice, there will not be a molecule with zero transverse motion in the decelerator, this one-dimensional calculation for the timing molecule is helpful for determining how an entire packet of molecules should be decelerated. The timing molecule is also referred to as the synchronous molecule, but in order to be consistent with other types of Stark deceleration that will be discussed later, it is called the timing molecule here. In practice, an entire molecular packet centered around the timing molecule will be decelerated.

As the potentials alternate, the potential well that the timing molecule occupies travels down the length of the decelerator. In the simplest method of running the decelerator, no deceleration occurs and the average velocity of the potential wells is the same as the initial velocity of the

beam. This configuration is called bunching because the molecular packet will not be decelerated, but instead will be forced into “bunches” of molecules. Bunching occurs if the potentials switch after a time $t = \frac{l}{v_i}$ where l is half the distance between Stark potential peaks in one configuration. Figure 2.4 shows the potential wells from Figure 2.3 overlaid in blue and red dashed curves, while the black curve gives the trajectory of the molecule on the two potential surfaces. Each time the molecule switches from one potential to another is referred to as one stage of deceleration. In this configuration, the molecule gains the same amount of energy that it loses during each stage.




Figure 2.4: The black curve depicts the potential experienced by the timing molecule as it travels through the two potentials (red and blue) for $\phi_0 = 0^\circ$, the case where the molecule does not change velocity.

The configurations are switched each time the molecule is at the exact same position on the longitudinal potential. The position of the timing molecule during the switch is typically given as an angle and is called the phase angle, ϕ_0 . The conversion between the position of the timing molecule within a well, z_0 , and the phase angle in radians is

$$\phi_0 = \frac{\pi z_0}{l}, \quad (2.1)$$

where $z_0 = 0$ is the position at which $\phi_0 = 0^\circ$. The phase angle gives a measurement of the longitudinal position within the decelerator that is independent of that pin separation, l , which

makes it both a useful, and sometimes frustrating, parameterization. The phase angle gives the longitudinal distance of the timing molecule as a fraction of the pin stage length, and has units of degrees or radians. It is a useful characterization of the magnitude of deceleration since it gives a sense of whether a lot or a little energy is removed per deceleration stage, but it does not give any information as to the final velocity of the beam.




Figure 2.5: The longitudinal potential for the given electrode configuration is given in red. Various phase angles, ϕ_0 , relative to this potential are marked.

Figure 2.5 shows where the phase angles of 0° , 45° , and 90° are located relative to a given potential (red). The phase angle $\phi_0 = 0^\circ$ is defined as the position of the timing molecule in the case where the final velocity of the molecule is the same as the initial velocity. Figure 2.4 shows the longitudinal trajectory of a molecule that switches potentials each time it reaches the position $\phi_0 = 0^\circ$. This molecule neither gains nor loses energy. $\phi_0 = 0^\circ$ is not at the bottom of a well but instead at a distance $l/2$ away from a pair of pins. The height of the Stark potential at $\phi_0 = 0^\circ$ will most likely not be at half the amplitude of the longitudinal potential. Instead the energy where $\phi_0 = 0^\circ$ occurs will depend on the gradient of the potential well. The phase angle increases from $\phi_0 = 0^\circ$ to $\phi_0 = 90^\circ$ for $z_0 = 0$ to $z_0 = l/2$. A phase angle of $-90^\circ < \phi_0 < 0^\circ$, which occurs at

$l/2 < z < l$, accelerates the molecule since the timing molecule will always be traveling down the potential energy hill. Phase angles of $90^\circ < \phi_0 < 270^\circ$ are unstable.

Figure 2.6 shows deceleration with phase angles of $\phi_0 = 30^\circ$ and 60° . Here, the molecule changes potential energy surfaces each time the voltages are switched as indicated by a dashed vertical line. As the phase angle increases from $\phi_0 = 30^\circ$ to $\phi_0 = 60^\circ$, more energy is removed per stage, indicated by ΔK . The deceleration extreme is a phase angle of $\phi_0 = 90^\circ$, which switches when the timing molecule reaches the peak of the potential energy hill. This is the maximum energy that can be removed per stage. The total kinetic energy removed from a timing molecule will be

$$K_{total}(\phi_0) = \Delta K(\phi_0) \times \text{number of stages.} \quad (2.2)$$

For typical pin decelerator geometry (>100 stages) and voltages, a small molecule such as OH or ND₃ will have $\Delta K(\phi_0 = 50^\circ) \approx 1$ K (or 1 cm^{-1}). However, this is only an estimate and ΔK is dependent on the Stark shift of the molecule being decelerated.

The longitudinal kinetic energy removed from a molecule by the entire decelerator and for a given phase angle (Equation 2.2) is dependent on electrode geometry, which sets ΔK and length of the decelerator. This can only be calculated if the gradient of the potential is known. Thus $K_{total}(\phi_0)$ is usually not a very intuitive function (ϕ_0) and it is difficult to guess the final velocity of the timing molecule from the phase angle. The relationship between the final velocity and phase angle for a timing molecule that undergoes n deceleration stages, each of which removes ΔK of energy, can be found via conservation of energy

$$\begin{aligned} \frac{1}{2}mv_i^2 &= \frac{1}{2}mv_f^2 + n\Delta K(\phi_0) \\ v_f(\phi_0) &= \sqrt{v_i^2 - \frac{2n\Delta K(\phi_0)}{m}}. \end{aligned} \quad (2.3)$$

The value of $\Delta K(\phi_0)$ is determined by building a model of the decelerator pins in *COMSOL Multiphysics*, a finite element analysis software package, and solving for the electric fields. While the full decelerator is many stages long, only a small section that is $2l$ long needs to be modeled, since the potentials repeat on this length scale.

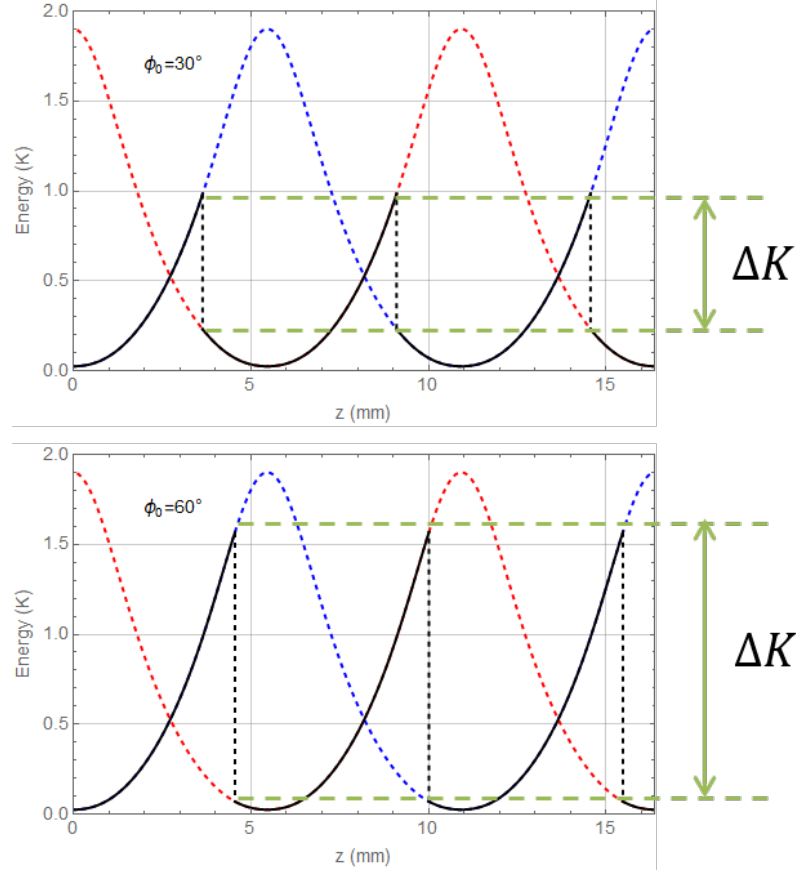


Figure 2.6: The two potentials that can be formed in PPSD are shown in blue and red. For each phase angle, $\phi_0 = 30^\circ$ and 60° , the path that the timing molecule travels is given in black. The times when the molecule reaches vertical-dashed lines indicate when the voltage configurations should switch.

To effectively implement deceleration at a given phase angle, a timing sequence for when the potentials switch must be calculated. The velocity of the timing molecule decreases with each stage, which increases the time each stage remains on as the molecule travels the length of the decelerator. The timing sequence for switching the potentials for a given phase angle is generated by simulating the motion of a timing molecule along the COMSOL generated longitudinal potential. The molecule is fixed to have no transverse displacement or transverse velocity, so its motion only on the longitudinal axis needs to be considered.

To find the time that the timing molecule spends in each stage, we use

$$v_f(\phi_0) = \frac{d\phi}{dt} = \frac{\pi}{l} \frac{dz}{dt} \quad (2.4)$$

and Equation 2.3 with $n = 1$, because we are referring only to one stage, to get

$$dt = \frac{l}{\pi} \frac{dz}{\sqrt{v_i^2 - \frac{2\Delta K(\phi_0)}{m}}},$$

$$t_{stage} = \frac{l}{\pi} \int_{-\pi+\phi_0}^{\phi_0} \frac{dz}{\sqrt{v_i^2 - \frac{2\Delta K(\phi_0)}{m}}}. \quad (2.5)$$

Here, v_i is now the initial velocity of the timing molecule at a given stage. The process of determining t_{stage} is usually completed numerically. Using the COMSOL model of the decelerator, an acceleration file is generated. It specifies the acceleration of a molecule along the central axis of the decelerator. A single timing molecule with an initial longitudinal z-velocity is propagated along z using a time step of ~ 10 ns. Each time the timing molecule reaches the switch point (phase angle), the acceleration file is shifted by l in the z-direction and the time recorded. The final list of switch times then corresponds to the switching sequence for that phase angle. The final velocity can also be extracted from this calculation, though it is a much more complex method than Equation 2.3 it gives identical results.

Up until now, we have considered only one timing molecule moving in the Stark potential well; in reality, an entire packet of molecules will be decelerated. Here, it is important to consider the packet of molecules in phase space, the 6D space of position and velocity. The area that successfully decelerated molecules occupy is known as the phase-space acceptance. The molecules that are near the timing molecule in phase space are the ones that will also be decelerated, and Figure 2.7 gives an example of this using $\phi_0 = 45^\circ$. The timing molecule on the black trajectory switches from the blue potential to the red potential when it reaches the phase angle $\phi_0 = 45^\circ$. The lagging molecule a (orange), which was behind the timing molecule on the blue potential, ends up at a higher Stark potential than the timing molecule on the red potential. This causes molecule a to lose less energy than the timing molecule. This allows the lagging molecule to move a little faster and to “bunch” with the timing molecule. The opposite is true of molecule b (green), which was ahead of the timing molecule. Molecule b experiences a greater energy loss when it switches to the red potential. This effect allows the rest of the packet to appear to catch up with the forward molecule. Only molecules

a given distance away in phase space from the timing molecule will experience this restoring force. Once a molecule is outside the phase-space acceptance, it cannot maintain a stable orbit inside the decelerating potential well and will not be decelerated. In the case of $\phi_0 = 90^\circ$, molecule b will have already crossed over to the accelerating portion of the potential well and all other molecules ahead of the timing molecule will not be decelerated. This model shows how phase angles close to $\phi_0 = 90^\circ$ are less effective for decelerating appreciable numbers of molecules.




Figure 2.7: Molecules traveling on the blue potential behind (orange) and ahead (green) of the timing molecule (black) experience longitudinal bunching due to the differing amounts energy loss when they switch to the subsequent (red) potential.

The alternating sets of crossed pins in the decelerator provide alternating transverse focusing. For energized stages, the highest electric field is near the rods, and on any particular stage, the molecules are focused toward the electric field minimum lying halfway between the two rods. By alternating the directions of the pins, there is overall transverse focusing.

There are three main effects that reduce the decelerator efficiency, one that can occur at any velocity and two that predominantly occur at velocities below 100 m/s. At any velocity in the decelerator, coupling between the longitudinal and transverse motions can amplify the transverse motion of molecules on each stage of deceleration, thus forcing the molecules transversely out of the decelerator [139]. At velocities <100 m/s in the decelerator, the phase-space acceptance is

drastically reduced due to “over aggressive” deceleration of portions of the packet through two effects, usually referred to as transverse over focusing and longitudinal reflection [140]. At low velocities of the decelerated packet, it becomes more likely for a stage to over focus in the transverse direction and squeeze the molecules out of the well (transverse over focusing). A second effect is that some low velocity molecules may not make it over the last deceleration stage when the average longitudinal kinetic energy of the packet is of the order of the potential energy hill (longitudinal reflection). Longitudinal reflection is particularly applicable to molecules away from the central axis where the longitudinal potential is higher. At these low final velocities, the molecules spend longer in the decelerator and have a greater chance to be lost.

2.2.2 Traveling-wave Stark Deceleration

The losses that are present in pin decelerators can be mitigated by using ring-shaped electrodes. The ring shaped electrodes can use a fundamentally different mode of applying voltages to the electrodes. Instead of pulsed voltages, chirped sinusoidally varying voltages are applied to the ring electrodes. This produces a traveling-wave potential well that continuously guides or decelerates the molecular beam. While PPSD produces a time averaged moving potential well, TWSD produces a genuine moving three-dimensional (3D) potential well. This provides an inherently stable deceleration process that has a true 3D trapping potentials and no longitudinal and transverse coupling loss. Continuous application of the decelerating force also allows for lower instantaneous accelerations and more efficient deceleration.

A ring decelerator is made up of a series of ring electrodes with every p rings electrically connected. Just like in PPSD, each set of rings that are electrically connected are referred to as a rod. The rods i are typically labeled $i = 0, \dots, p - 1$. Figure 2.8 shows a CAD drawing of one ring (a) and our ring decelerator (b). In this case, the rings are tennis racket shaped with the two wires flaring out at the end. Other ring decelerators [122] have not use the flared configuration, and instead kept the two wires together to equal success. The eight ring periodicity ($p = 8$) of the decelerator can be seen in Figure 2.8b, where a CAD drawing has each rod, and its corresponding




Figure 2.8: CAD drawings of one ring electrode (left) and the full decelerator (right). Each ring is held in place on its mounting rod with two set screws and every eighth ring is electrically connected. The decelerator has been color coded to more easily distinguish the periodicity. Even though the ring is not a perfect circle, the rings are mounted in a spiraling pattern and the pinch does not significantly contribute to any changes in the potential well from a perfect ring [95]. Taken from [95] Figure 5.16

rings, color coded. Figure 2.1b shows the end of the decelerator in our lab.

Instead of each rod being on or off at a given time, in TWSD, the ring electrodes have varying sine-wave voltages $V_i(t)$ on each rod i . The timing molecule no longer climbs successive potential energy hills, instead it travels along with the potential well. In order to create this moving well, the sine-wave voltages on the rings will have the form

$$V_i(t) = V_a \sin\left(\frac{2\pi i}{p} - \omega t\right) \text{ for } i = 0, \dots, p - 1. \quad (2.6)$$

Here ω sets how fast the potential well translates in time. If $\omega = 0$, then there is only a static potential well. If $\omega \neq 0$, a constant frequency longitudinally translating sine-wave voltage, which would correspond with bunching, would be produced. For a decelerator with ring-to-ring spacings of l , the velocity of the well is the distance the well travels in one period, pl , multiplied by the frequency f of the sine wave, $v = fpl$. For bunching at the initial velocity of the beam, $\omega = \frac{2\pi v_i}{pl}$.

In order to decelerate a beam, the velocity of the well itself must change. This is accomplished

by chirping the frequency of $V_i(t)$. The simplest frequency chirp is a linear one where

$$\begin{aligned} f(t) &= f_0 + \alpha t, \\ \omega(t) &= \omega_0 + 2\pi\alpha t, \end{aligned} \quad (2.7)$$

and where α is a constant acceleration given by

$$\alpha = \frac{f_f - f_0}{T} = \frac{v_f - v_i}{plT}. \quad (2.8)$$

Here T is the total time required to travel the length of the decelerator and is determined by

$$T = \frac{2L}{v_f + v_i}. \quad (2.9)$$

L is the full length of the decelerator and l is the distance between adjacent rings; they are related by

$$L = l \times (\text{number of rings} - 1). \quad (2.10)$$

In order to incorporate the frequency chirp, now instead of having a ωt term, $V_i(t)$ must have a time dependent phase $\Phi(t)t$ term, and Equation 2.6 becomes

$$V_i(t) = V_a \sin\left(\frac{2\pi i}{p} - \Phi(t)t\right) \text{ for } i = 0, \dots, p-1. \quad (2.11)$$

$\Phi(t)$ is the integral of the instantaneous frequency (Equation 2.7),

$$\Phi(t) = \int \omega(t) dt = \omega_0 t + \pi\alpha t^2. \quad (2.12)$$

Now by substituting Equations 2.8 and 2.9 into Equation 2.12, $\Phi(t)$ can be reexpressed as

$$\begin{aligned} \Phi(t) &= \omega_0 t + \frac{\pi}{2plL} (v_f^2 - v_i^2) t^2 \\ \Phi(t) &= \frac{2\pi v_0}{pl} t - \frac{\pi}{2plL} (v_i^2 - v_f^2) t^2 \\ \Phi(t) &= \frac{2\pi}{pl} \left[v_0 - \frac{v_i^2 - v_f^2}{4L} t \right] t. \end{aligned} \quad (2.13)$$

The full chirped sine-wave voltage (Equation 2.11) then becomes

$$V_i = V_a \sin\left(\frac{2\pi i}{p} - \frac{2\pi}{pl} \left[v_0 t - \frac{v_f^2 - v_i^2}{4L} t^2 \right]\right) \text{ for } i = 0, \dots, p-1. \quad (2.14)$$




Figure 2.9: The 3D (A) and cross section (B) view of the instantaneous electric potentials for ± 12 kV sine-wave voltages on the ring electrodes and (C) the resulting molecular potentials in kelvin for ND_3 . For a sine-wave potential spanning eight rings, two Stark potential wells are formed. Each well is 0.7 K deep along the longitudinal axis. Figure taken from [95] Figure 5.1.

These voltages create 3D moving potential wells on the ring electrodes. Figure 2.9 shows the electric potential created by sine-wave voltages on the ring electrodes and the resulting molecular potential wells for ND_3 . Although one period of the sine-wave covers eight rings, two potential wells are created for every eight rings because the Stark energy depends only on the magnitude of

the electric field and not the sign (see Section 3.1.1). For a maximum voltage of ± 12 kV on the ring electrodes, each well is 0.7 K deep along the longitudinal axis. While this is shallower than typical wells in PPSD (~ 2 K), continuous confinement of molecules in TWSD yields more efficient deceleration. The instantaneous accelerations applied to the molecules is lower, but it is applied constantly instead of discretely. The cylindrical symmetry and continuous nature of TWSD makes it a more efficient deceleration technique than PPSD, but the required sine-wave voltages make it more difficult to implement.

2.2.3 Pulsed Ring Stark Deceleration

Since TWSD is difficult to implement, pulsed-ring Stark deceleration, an alternative running mode for operating a ring-geometry decelerator was proposed by Hou *et al.* [141]. It uses pulsed voltages, just like in PPSD, but on ring electrodes. This creates discrete potential wells similar to PPSD, but with a slightly different potential well shape. The voltages on the electrodes are varied in an identical manner to that of PPSD. This makes many of the techniques from PPSD, such as calculating the timing sequences and the phase-space acceptance, directly transferable between the two methods. PRSD takes advantage of the cylindrical symmetry provided by the rings and is simpler to implement than TWSD, since it uses commercial high-voltage switches instead of high-voltage analog amplifiers. The PRSD method is discussed in detail in Chapter 4.

2.3 Phase-space Acceptance

A molecule that travels only along the axis of the decelerator would be stable in an infinitely long decelerator. However, in reality, the molecules have off-axis velocity and position components. In all deceleration modes, the transverse potentials within the decelerator drive these molecules towards the decelerator axis. At the same time, the longitudinal potentials push molecules ahead of and behind the timing molecule toward the timing molecule (see Section 2.2.1).

Depending on the geometry of the decelerator, voltage on the electrodes, and the phase angle, only molecules a given distance away in position and velocity from the timing molecule will

be successfully decelerated. The boundary in phase space between the molecules that are stably decelerated and those that are not is called the separatrix. The area inside the separatrix is called the phase-space acceptance. If considering only the position and velocity along one coordinate (say z and v_z in the case of the longitudinal phase-space acceptance) this is a 2D phase-space acceptance. However, in the case of all six phase-space coordinates, the phase-space acceptance is 6D and has units of $(\text{mm})^3 (\text{m/s})^3$.

Ideally, the phase-space acceptance will be well filled with molecules from the initial molecular pulse that enters the decelerator, however this is not always the case. This incomplete filling can be due to two reasons, the first is instabilities in the deceleration process that cause holes within the phase-space acceptance and the second is the initial pulse was smaller in phase-space than the phase-space acceptance of the decelerator. The first issue will play a role in PPSD since it has well-characterized instabilities and loss mechanisms that are enhanced at low final velocities [139]. In particular, coupling between the longitudinal and transverse motion leads to amplification of the transverse motion. The second issue of underfilled phase space by the initial molecular packet will not be discussed because the initial pulse used in these experiments is much larger in phase space than the acceptance of either decelerator.

The phase-space acceptance of the decelerator can be calculated via two methods, by examining the restoring force on molecules near the timing molecule and calculating their equations of motion, or by looking at the combination of the fictitious force from the deceleration of the well and the Stark potential well. These two methods are considered in 2D slices in the longitudinal (z) directions. Here, we also address only S=1 deceleration, where a set and uniform amount of energy is removed in each stage of the decelerator.

In the first method, we consider the motion of the timing molecule and the difference in the force applied to the timing molecule and a molecule near the timing molecule. The separatrix is then the boundary between stable and non-stable orbits of these molecules. First, we examine the timing molecule decelerated with a phase angle of $\phi_0 = \frac{\pi z_0}{l}$ (Equation 2.1). To find the average force on the timing molecule per deceleration stage, an approximate functional form for the potential

well is required. The longitudinal potential hill can be expressed as a Fourier series of the form

$$W(z) = \frac{1}{2}a_0 + \sum_{n=1}^{\infty} a_n \cos(nz) + \sum_{n=1}^{\infty} b_n \sin(nz). \quad (2.15)$$

The potential should be expressed as a function of the longitudinal distance (ϕ), which is given in a similar fashion to the phase angle (ϕ_0) such that $\phi = \frac{\pi z}{l}$. We change the potential $W(z)$ from the longitudinal distance z to the phase angle ϕ and expand around $\phi_0 = \frac{\pi}{2}$. a_0 is a constant energy shift and is set to zero for now. All b_n terms vanish since the potential is symmetric about ϕ_0 . If the rods are spaced such that l , the longitudinal spacing between adjacent rods, is approximately equal to the transverse center-to-center spacing, the potential closely resembles just the first term of a_n , a cosine wave [142]. This gives the potential $W(\phi)$ as

$$W(\phi) \approx a_1 \cos\left(\phi + \frac{\pi}{2}\right). \quad (2.16)$$

The coefficient a_1 is the amplitude of the cosine wave potential hill. Thus, if the height of the potential well is W_{max} then $2a_1 = W_{max}$.

The average force on the timing molecule per deceleration stage at phase angle ϕ_0 can be determined from the loss of kinetic energy per stage. This is illustrated in Figure 2.10 where $\Delta K = E_{red} - E_{blue}$. Since $E_{blue} = E_{green}$ and the energy at E_{green} is $W(\phi_0 + \pi)$,

$$\begin{aligned} \bar{F}_0 &= \frac{\Delta K}{l} \\ \bar{F}_0 &= \frac{1}{l} (W(\phi_0) - W(\phi_0 + \pi)). \end{aligned} \quad (2.17)$$

A substitution of Equation 2.16 into Equation 2.17 gives an average force of

$$\begin{aligned} \bar{F}_0 &= \frac{a_1}{l} \left(\cos\left(\phi_0 + \frac{\pi}{2}\right) + \cos\left(\phi_0 + \frac{3\pi}{2}\right) \right) \\ \bar{F}_0 &= \frac{2a_1}{l} \sin(\phi_0) \\ \bar{F}_0 &= \frac{W_{max}}{l} \sin(\phi_0). \end{aligned} \quad (2.18)$$

Any molecules near the timing molecule will experience a similar force, but off by $\Delta\phi$, the longitudinal distance it is away from the timing molecule in radians. The average force on a molecule

located at $\phi_0 + \Delta\phi$ is

$$\bar{F} = \frac{W_{max}}{l} \sin(\phi_0 + \Delta\phi), \quad (2.19)$$

and the average force difference between this molecule and the timing molecule is

$$\bar{F}_{diff} = \bar{F}_0 - \bar{F} = \frac{W_{max}}{l} [\sin(\phi_0) - \sin(\phi_0 + \Delta\phi)]. \quad (2.20)$$




Figure 2.10: The timing molecule switches from the red potential to the blue potential at the phase angle ϕ_0 . The energy lost per stage is ΔK and is the same as $E_{red} - E_{blue} = E_{red} - E_{green} = W(\phi_0) - W(\phi_0 + \pi)$.

Now, the equation of motion for a molecule located near the timing molecule and relative to the timing molecule can be written down from Newton's second law,

$$\begin{aligned} m \frac{d^2 z}{dt^2} &= F_{diff} \\ \frac{ml}{\pi} \frac{d^2 \phi}{dt^2} &= F_{diff} \\ \frac{ml}{\pi} \frac{d^2 \phi}{dt^2} &= \frac{W_{max}}{l} [\sin(\phi_0) - \sin(\phi_0 + \Delta\phi)] \\ 0 &= \frac{ml}{\pi} \frac{d^2 \phi}{dt^2} - \frac{W_{max}}{l} [\sin(\phi_0) - \sin(\phi_0 + \Delta\phi)]. \end{aligned} \quad (2.21)$$

Molecular trajectories can now be calculated relative to the timing molecule. The boundary between stable and unstable trajectories around the timing molecule is the separatrix. In order to calculate the phase-space acceptance, it would be easiest to work in z and V_z where the timing molecule is defined to be at $z = 0$ m and $V_z = 0$ m/s. We use $\Delta\phi = \frac{\pi z}{l}$ to change Equation 2.21 to z . Although $\phi_0 = \frac{\pi z_0}{l}$, ϕ_0 is left as is since ϕ_0 and z_0 are constants and pulsed deceleration is typically described by the phase angle in degrees and not the exact distance, z_0 , that the timing molecules moves.

Equation 2.21 now becomes

$$0 = \frac{d^2 z}{dt^2} - \frac{W_{max}}{ml} \left[\sin(\phi_0) - \sin\left(\phi_0 + \frac{\pi z}{l}\right) \right]. \quad (2.22)$$

Equation 2.22 can be numerically integrated to determine the separatrix for a given phase angle ϕ_0 . Figure 2.11 gives an example of a few separatrices obtained for PPSD of ND₃. The area bound by the a given separatrix is the phase-space acceptance.




Figure 2.11: The longitudinal (z axis) separatrices for $\phi_0 = 0^\circ, 20^\circ, 40^\circ, 60^\circ, 80^\circ$. The area bound by a given separatrix is the phase space acceptance. The z coordinate here is the distance away from the position of the timing molecule.

The second method of determining the phase-space acceptance examines the non-inertial reference frame of the decelerating potential relative to the inertial lab frame. In the non-inertial frame, the molecules experience an opposing fictitious force F_f

$$\begin{aligned}
 -\frac{dU_f}{dz} &= F_f \\
 -\frac{dU_f}{dz} &= -ma \\
 U_f &= \int ma \, dz \\
 U_f &= maz,
 \end{aligned} \tag{2.23}$$

where a is the average acceleration of the potential well and U_f is the fictitious potential. Since we are discussing deceleration, a is negative. The effective potential that the molecules experience is the sum of the actual potential, U , and the fictitious potential

$$U_{eff} = U + maz. \tag{2.24}$$

In a pulsed mode of deceleration, the time-averaged effective potential, U_{eff} , can be found from the force difference (Equation 2.20) between a molecule located at z and the timing molecule,

$$\begin{aligned}
\bar{F}_{diff} &= \frac{W_{max}}{l} [\sin(\phi_0) - \sin(\phi_0 + \Delta\phi)] \\
-\bar{F}_{diff} &= \frac{W_{max}}{l} \left[\sin\left(\phi_0 + \frac{\pi z}{l}\right) - \sin(\phi_0) \right] \\
\frac{d\bar{U}_{diff}}{dz} &= \frac{W_{max}}{l} \left[\sin\left(\phi_0 + \frac{\pi z}{l}\right) - \sin(\phi_0) \right] \\
\bar{U}_{diff} &= \int \frac{W_{max}}{l} \left[\sin\left(\phi_0 + \frac{\pi z}{l}\right) - \sin(\phi_0) \right] dz \\
\bar{U}_{diff} &= -m \left[\frac{W_{max}}{ml} \sin(\phi_0) \right] z + \frac{W_{max}}{\pi} \cos\left(\frac{\pi z}{l} + \phi_0\right). \tag{2.25}
\end{aligned}$$

Again, here, we use the slightly strange convention of $\phi_0 + \frac{\pi z}{l}$ which has units of radians. Equation 2.25 has been written in a suggestive format that matches Equation 2.24. The term inside the square brackets is the relative acceleration

$$a = \frac{W_{max}}{ml} \sin(\phi_0). \tag{2.26}$$

By multiplying Equation 2.26 by m we get \bar{F} or Equation 2.19 for the timing molecule. The far right term of Equation 2.25 is the averaged potential well centered on the timing molecule. It is a factor of $\frac{2}{\pi}$ different from Equation 2.16 due to the need for averaging across one deceleration state in the pulsed mode.

In the traveling-wave mode, since the forces are applied continuously and the potential does not change significantly relative to the timing molecule, Equation 2.24 does not require the time-averaged force derived in Equation 2.25. Instead U is simply the Stark potential well (usually calculated via COMSOL) and the acceleration is

$$a = \frac{v_f^2 - v_i^2}{2L}, \tag{2.27}$$

where L is the length of the decelerator.

This second method allows for easy visualization of the possible phase-space acceptance without having to do much calculation. By plotting Equation 2.24, it is easy to see when the Stark potential well ‘tilts’ over to the point of no longer being able to trap molecules. Figure 2.12 shows

several longitudinal potentials for ± 8 kV TWSD in the ring decelerator in our lab. The initial velocity is set to 415 m/s. If no deceleration is applied, the well shape remains symmetric. As more acceleration is applied the well tilts and the downstream (right) wall of the potential well lowers in height. For this given potential well, deceleration down to 100 or 0 m/s is not very feasible since the potential well almost completely flattens out.

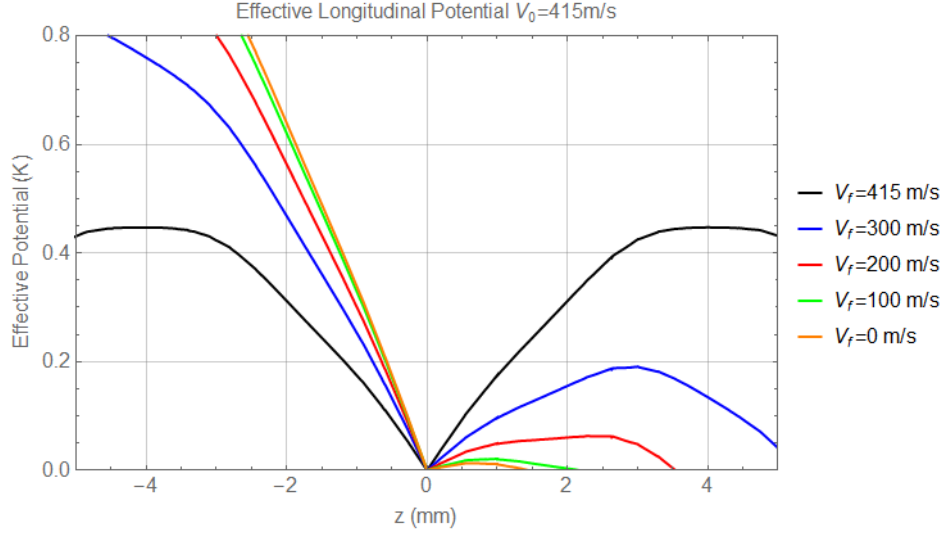


Figure 2.12: The effective longitudinal potential (Equation 2.24) for an initial velocity of 415 m/s and various final velocities for ± 8 kV TWSD. As more acceleration is required to reach lower final velocities, the well tilts over significantly. For this voltage configuration, final velocities of 100 and 0 m/s have almost no potential wall on the right-hand-side. Thus very few molecules can be phase stable in that deceleration sequence.

The longitudinal phase-space acceptance can also be calculated using the U_{eff} potential (Equation 2.24) and conservation of energy. The separatrix is set by molecules that can just escape the U_{eff} potential well. In order to stay phase stable, molecules must have less than or equal energy to the downstream (lower) maximum Stark energy,

$$\frac{1}{2}mv^2 + U_{eff} = U_{eff_{max}}. \quad (2.28)$$

Solutions to Equations 2.22 and 2.28 produce identical phase space acceptances.

Up to this point, all pulsed deceleration phase spaces have been examined using the first term in a fourier series as the potential. While this is an acceptable approximation, it is important to

note that there are differences in the shape of the potential. Figure 2.13 shows the longitudinal potentials for all three types of deceleration (TWSD, PRSD, and PPSD) in addition to a the cosine potential given by Equation 2.16. The cosine potential more closely resembles the potential well in the pin decelerator than the ring decelerator. In the ring decelerator the vertex at the bottom of the well cannot be easily modeled with a cosine. The difference between the cosine potential and actual potential will change the shape of the separatrix and also affect the acceleration as a function of phase angle/final velocity.

In the cosine wave approximation of the potential, the acceleration is given by Equation 2.26. However, in the actual experiment, the average acceleration is determined by Equation 2.27. Although the acceleration is not applied in the same fashion in pulsed and traveling-wave deceleration, once averaged over the total deceleration time, they are the same. This is only true because we use a constant acceleration in all our experiments. If this was changed, different calculation would be required. Figure 2.14 shows the acceleration as a function of phase angle for both equations, where black is the calculation using Equation 2.27 and blue is the calculation using Equation 2.26. Red is a sanity check and is the acceleration averaged over one stage of deceleration. It is determined by taking average value of the piece-wise acceleration function at different phase angles. The acceleration determined using this method matches Equation 2.26 exactly.




Figure 2.13: Normalized longitudinal potentials for TWSD, PRSD, PPSD, and a cosine (Equation 2.16). All potentials have been normalized so that they have a height of one and the same longitudinal periodicity. Both potentials in the ring decelerator the have a vertex at the bottom of the potential well.




Figure 2.14: The acceleration for the PRSD potential as a function of phase angle. The acceleration calculated from v_i and v_f (black) and from averaging over one deceleration stage (red) are identical, so the black curve appears underneath the red curve.

2.4 Simulations

In order to fully explore the dynamics in any deceleration mode, three-dimensional (3D) molecular trajectory simulation are required. These 3D simulations are particularly important for the PPSD mode, due to coupling effects that cannot be replicated in a 1D model. These simulations are semiclassical. The potential energy surfaces are determined by the Stark energy of

Table 2.1: Typical values for the full width at half maximum (FWHM) of the initial distribution of molecules

Phase-space value	Gaussian FWHM
x,y-position	3 mm
z-position	30 mm
x,y-velocity	5 m/s
z-velocity	30 m/s

the molecules (a quantum calculation), but the trajectories are treated classically. In general, the simulations follow this procedure:

- (1) Create a Gaussian distribution of initial position and velocities using Monte Carlo techniques. Table 2.1 gives typical initial parameters. The distribution is centered at the location of the pulsed valve.
- (2) One molecule, typically the first molecule, is set to be the timing molecule. It will have no transverse displacement or velocity and will be constrained to evolve only in the longitudinal direction. This molecule will determine when the potentials in the decelerator change.
- (3) Propagate the distribution to the entrance of the decelerator. This is the free flight section and no forces are exerted on the molecules.
- (4) Propagate the distribution through the decelerator via the chosen mode of deceleration. The potentials will change as dictated by the timing molecule and desired phase angle or final velocity.
- (5) Record the final phase-space distribution (a set time) or the time-of-flight (a set position) for the molecules.

Typically, modified Euler integration is used for the molecule trajectories in the simulations. The time scale (3-6 ms) for deceleration is short enough that no appreciable errors build up. For trap simulations, a more complex symplectic integrator must be used because errors can accumulate when holding molecules in the trap for an extended period of time (1-2 s).

In all simulations, acceleration files for the given molecule of interest are generated from electric potentials calculated in COMSOL. In pulsed mode simulations, the acceleration file is specified for a length of $2l$ (spanning three pairs of electrodes). This file is then propagated forward by l each time the decelerator potentials switch. In traveling-wave mode simulations, 25 versions of the accelerations are generated, which correspond to the bottom of the potential well being in 25 equally spaced locations between two rings. Again the potentials are propagated forward by l when the bottom of the potential well moves from ring n to ring $n + 1$.

Molecules may be lost from the simulation because they end up in the following locations: outside the entrance area of the decelerator, hitting a ring, or outside decelerator phase space. Molecules in the last category will typically travel the length of the decelerator, but not be decelerated. They will exit before the decelerated packet. They are not within the phase-space acceptance and are not decelerated. In the case of mild deceleration ($< |25| \text{ km/s}^2$), there could be many of these molecules since the undecelerated beam may not be fully separated out from the decelerated packet and not all molecules that arrive in the time window of the decelerated peak will be within the phase-space acceptance of the decelerator. The decelerated molecules can be selected for by plotting all molecules in phase space and picking only those that fall within the calculated 1D phase-space acceptance in each direction (x , y , and z).

Simulations are also able to provide information regarding what is happening inside the decelerator that can never be extracted from an experiment. These include the phase-space information for molecules at intermediate times and individual particle trajectories in the decelerator. To save on memory, the molecules that “crash” can be removed from the simulation, and recording when and where these molecules are lost can also be very helpful for studying various deceleration schemes.

These simulations are written in C++ and run on the JILA’s ‘Terra’ cluster. Depending on the initial density of molecules and the number of desired decelerated molecules, it is often desired to begin a simulation with several million molecules. Instead of propagating all the molecules at once, since there are no inter-molecule dependencies in the simulation, many smaller distributions

of molecules can be run in parallel or series. All results can be summed at the end.

In all of these simulations, the timing molecule is artificial and must be removed from the molecular distribution before the results are analyzed. This is especially true if running many iterations or very spread out initial packets. For example, if 100 iterations of 10,000 molecules were used to generate 1 million molecules. One hundred of the molecules were forced to arrive in the decelerated packet. If the decelerated packet is small, only a few hundred molecules, then the timing molecules will artificially boost the decelerated and detected molecules.

Trapped molecule simulations are conducted in a similar fashion to the process outlined above. The Stark potentials created by the electrostatic trap are modeled in COMSOL and exported. The molecule distribution is taken from the output of a trap loading simulation. Since time spent in the trap is significantly longer than time spent in deceleration or trap loading (seconds compared to milliseconds), it is critical to use a symplectic integrator. The total energy of individual molecules in the trap should be checked to ensure that they neither gain nor lose significant energy.

2.5 Experimental Layout

Figure 2.15 gives the layout for a typical deceleration experiment that contains a pulsed valve, skimmer, decelerator, interaction and detection region. The configuration of the interaction and detection region may vary widely depending on the intended experiment.

The molecule of interest is seeded in a carrier gas, and the gas mixture enters the vacuum chamber via a pulsed valve. The supersonic expansion of the gas from the gas reservoir into the vacuum system cools the internal degrees of freedom of the molecule, but the beam maintains a large forward velocity. Typically, krypton or xenon is used as the carrier gas in order to have a low initial velocity even though heavier carrier gases are also more likely to form clusters, which cause higher temperatures.

The beam passes through a skimmer, which both selects for a collimated beam and allows for differential pumping between the source chamber and decelerator chamber. This way the majority of the gas loaded into the vacuum system remains in the source chamber. The flared geometry of




Figure 2.15: A general Stark deceleration experiment contains a pulsed valve, skimmer, decelerator, and an interaction and detection region. Here a pin-geometry decelerator is depicted and the detection system is a 1+1' resonance-enhanced multiphoton ionization (REMPI), time-of-flight (TOFMS), and microchannel plate (MCP) system.

the skimmer reduces interference from molecules reflected off of the outside wall and build up of molecules inside the skimmer. Other versions of a skimmer, including a cooled skimmer, have shown greater cold molecule densities in the beam, but have not been tested with a decelerator [143].

After the skimmer is the Stark decelerator, which can come in several different geometries and electronic configurations as discussed in Section 2.2. The decelerator should be placed as close as possible to the skimmer to reduce the molecules lost in free flight after the skimmer and before the decelerator. While the pressure in the source chamber is on the order of 10^{-5} torr, to conduct trapped molecule experiments with satisfactory trap lifetimes, the pressure in the science chamber should be under 1×10^{-9} torr. To achieve these pressures, additional differential pumping will most likely be required in the decelerator chamber. This can be achieved by adding separation plates between pumping sections. The plates should closely surround, but not touch, the decelerator.

At the end of the decelerator, the decelerated molecules can be trapped, guided, or collided with other species. After the experiment has been completed, the molecules are typically detected in a state-selective manner. Molecules are usually detected at the end of the decelerator, but they can also be detected in the middle of the decelerator [118].

2.5.1 Supersonic Expansion

Most Stark deceleration experiments begin with a supersonic expansion from a pulsed valve, however they could also be loaded from effusive [73] or buffer gas beams [43]. Pulsed valves come in many varieties and operate using one of three mechanisms: forces on current carrying conductors, solenoids, or piezoelectric actuators (PZT). Current-loop valves use pulses of opposing currents on two parallel conductors to force apart the two conductors, which opens an o-ring seal [144,145]. One current loop valve worth noting is the Jordan *Pulsed Supersonic Valve C-211*, which is commercially available. Another valve that operates on current carrying conductors is the Nijmegen pulsed valve [146], which can be built at a fairly low costs. In this valve, a current is run through a metal strip placed between two magnets. The geometry of the system causes the metal strip to lift when a current is applied. The second type of valve, a solenoid valve, uses a solenoid to actuate a ferromagnetic arm connected to a poppet, which seals the valve opening [147]. These valves include the *Parker Series 9* and *Series 99*, which were originally produced by General Corp and are often referred to as *General Valves*, and the Even-Lavie valve [148]. The last design is the PZT valve [149,150]. The PZT flexes when a voltage is applied across it, retracting a poppet, which has an o-ring at the end. When the poppet is pulled back, in both the solenoid and PZT valves, a pulse of gas is let into the chamber.

Stark deceleration is a conservative process, meaning no cooling or phase-space compression will occur during deceleration. Gas expansion through the valve opening sets the maximum density of the decelerated beam, which will set the maximum achievable signal for the experiment. This makes the choice of the molecular source very important. A slower beam is easier to decelerate, however it tends to have lower initial density and is more expensive to create.

Supersonic beams have a narrow velocity spread around a mean velocity, which makes them a good source of internally cold, but fast moving molecules. As the gas expands from the high-pressure gas reservoir through the aperture (typically ~ 1 mm) and into the vacuum region, it undergoes many collisions. These collisions cool the internal (vibrational and rotational) degrees of

freedom of the molecule of interest. In the reservoir, the molecules have no net flow, but once they expand into the vacuum chamber, they form a directed jet perpendicular to the valve aperture.

The forward velocity of the beam in the ideal case where all enthalpy of the gas is converted into kinetic energy is [151]

$$v(T_0) = \sqrt{\frac{2k_B T_0}{m} \frac{\gamma}{\gamma - 1}}, \quad (2.29)$$

where T_0 is the temperature of the gas reservoir, m is the mass of the atom or molecule, and γ is the adiabatic constant. For a monatomic ideal gas, $\gamma = \frac{5}{3}$. In a real system, the heat capacity ratio γ depends on the gas density and is a function of the pressure and temperature. $\gamma(p_0, T_0)$ should be used instead and calculated from the heat capacities at constant pressure, $C_p(p_0, T_0)$ and constant volume $C_v(p_0, T_0)$.

Lighter carrier gases produce denser beams at a faster velocity. However, this requires more aggressive deceleration in order to reach low velocities, which decreases the phase-space acceptance of the decelerator. Haas *et al.* [116] found that the trade-off between increased density and increased phase angle leads to similar decelerated signals. This means if more aggressive deceleration is possible, it is better to run with a lighter carrier gas since it will be cheaper. If more aggressive deceleration is not possible, either higher voltages need to be applied to the decelerator or the decelerator must be made longer.

From Equation 2.29, using a room temperature of 300 K, the velocities of an expansion of argon, krypton, and xenon give beam velocities of 559 m/s, 386 m/s, and 308 m/s, respectively. In practice, when working with a seeded gas, we observe slightly higher average velocities. For example, when working with ND₃ or OH seeded in krypton we set the timing sequence for a molecule with an initial velocity of 410-450 m/s.

Not all molecules that are typically Stark decelerated come ready in a bottle like ammonia; some must be created in the chamber by an electric discharge or photodissociation. For example, we create our second most commonly used molecule, OH, via a electric discharge of H₂O. A molecular pulse created from a electric discharge or dissociation will be not be as broad as one from a bottle,

since the time scale for the discharge pulse or dissociation laser pulse is short compared to the length of the molecular pulse and these types of sources will fill fewer wells in the decelerator than a molecule from a bottle.

Many factors must be taken into account when choosing the carrier gas and valve for a Stark deceleration experiment. The molecule and decelerator geometry set the maximum acceleration that can be achieved. Limits to the maximum acceleration may set a maximum on initial velocity, which in turn sets a which carrier gas can be used. The final decision on the valve will depend on the carrier gas, operational temperature, and any need for additional attachments such as discharge plates.

2.5.2 Detection

The two most common state-selective molecule detection methods are laser-induced fluorescence (LIF) and resonance-enhanced multiphoton ionization (REMPI). A simplified model of both detection schemes is shown in Figure 2.16. In LIF, the molecule is excited to an excited state and spontaneous emission from the excited state is detected, typically with a photomultiplier tube (PMT). In REMPI, the molecules are ionized in a step-wise resonant process. One or more photons excite the molecule to an intermediate state where it is then ionized. The ions are typically detected with a microchannel plate (MCP).

An $m + 1$ REMPI scheme means that m photons are required to reach the excited state and then an additional photon is required to ionize it. If the photons are not all of the same wavelength, then the second photon is denoted by $1'$ indicating that it is two-color REMPI process. Figure 2.16 depicts a 1+1 REMPI scheme where both photons are of the same wavelength and each step requires only one photon. Once ionized, the molecules are detected via a time-of-flight mass spectrometer (TOFMS). An electric field is applied that accelerates the ions toward a MCP detector. The arrival time at the MCP of the molecules is dependent on their mass-to-charge ratio.

Both types of detection processes are quantum-state selective since the first step to the excited state is tuned to a particular transition. If a molecule is not in that desired state, it will not be




Figure 2.16: A simple model of both LIF and 1+1 REMPI detection. In both cases a resonant photon excites the molecule from the $|1\rangle$ state to the $|4\rangle$ state. In LIF, the $|4\rangle$ state decays to another state which could be $|3\rangle$, $|2\rangle$, or $|1\rangle$ and emits a photon, this photon is the detected signal. Here the decay is depicted on the $|4\rangle \rightarrow |3\rangle$ transition. In REMPI, a second photon ionizes the molecule from the $|4\rangle$ state into the continuum. The resulting ion is detected. If the molecule started in the $|2\rangle$ state instead of the $|1\rangle$ state it does not reach an excited intermediate state, and molecules in the $|2\rangle$ state are not detected.

detected. This is opposed to single photon ionization processes, which are not state selective. Both LIF and REMPI detection tend to use wavelengths that are more easily accessible than one-photon ionization process in the UV. For LIF and REMPI processes, the first photon is typically in the near UV and can be produced by a tunable dye laser. Since both of these processes are state selective, only one quantum state can be detected at a time. However, the tunability of the first wavelength means that by adjusting this wavelength various states can be detected [152].

There are advantages to each technique. In REMPI, only a small portion of a trapped sample is ionized and detected while in LIF the entire sample can be detected at once. In REMPI where $m > 1$, the laser beam(s) are focused down, and only a small portion of the sample is detected in each run of the experiment. If the focus of the laser beam(s) can be scanned, then the shape of

the trapped molecules can be mapped out. Individual ions can be detected in REMPI making it a more sensitive system, which can be used to map out the density distribution of a trapped sample [153]. LIF uses a collimated beam to illuminate the entire trapped sample. The signal is limited by the accessible solid angle for collecting photons. This can be difficult, particularly in electrostatic traps, since only a small solid angle is available to collect photons and there are many shiny electrodes for the excitation photons to reflect off of.

Often the detection technique chosen is dependent on the molecule in question, the available lasers, and the dynamics of the study. Other detection techniques aside from LIF and REMPI have been used, including non-resonant multiphoton ionization followed by mass and velocity selection. This was first demonstrated with CH_3F [130]. This non-resonant method is not dependent on having a convenient and accessible excited state, and could open up the possibility of decelerating and detecting more molecules that do not have a convenient REMPI or LIF detection scheme. However, it more difficult to set up since it requires careful ion optics and a velocity map imaging system.

2.5.3 Our experiments

In our lab, we have two decelerators, a pin decelerator and ring decelerator. Besides the decelerators, both experiments are fairly similar. They used the same type of valve and detection system. The detection system must be adjusted for the molecule of interest, but both system can be used with both ND_3 and OH . The molecular structures and detection schemes of ND_3 and OH are discussed in more detail in Sections 3.1 and 6.2 respectively. In these experiments, we use primarily krypton as the carrier gas. Beams seeded in argon are too fast to decelerate to rest in either of our systems given the length of the decelerator, however argon is occasionally used to debug a free flight beam. The molecule of interest is mixed with krypton in a 1-2% mixture with a partial pressure of 40-60 psi of krypton.

The gas is let into the chamber via a PZT valve, which is typically operated with a 100 μs , 300-500 V pulse. Figure 2.17 shows the ND_3 time-of-flight distribution from our PZT valve in addition

to that of an in-house built NPV. The NPV has not been used with deceleration experiments, but it appears that it would produce a similar beam profile to that of the PZT valve. Both beams were detected after a skimmer and with 2+1 REMPI. A *Beam Dynamics Inc.* skimmer with a 2 mm orifice diameter is placed ~ 15 cm down stream of the valve face. The exact distance between the valve and skimmer is adjustable since the valve is mounted on a xyz -translation stage and its position can be adjusted from outside the vacuum chamber. The arm of the translation stage must be as straight as possible, so that the valve face is parallel to the entrance of the decelerator. Having the valve closer to the skimmer means a denser beam, however too close to the skimmer and the skimmer will cause significant beam interference.




Figure 2.17: Time-of-flight measurements on ND_3 for both the PZT and NVP valves. This measurement using a skimmed beam and detected ND_3 via 2+1 REMPI. Taken from [154]

The decelerator is placed 3 mm after the downstream end of the skimmer. In this dissertation, only REMPI is used, but both decelerator systems are built to also accommodate LIF detection. The laser for either method enters the vacuum chamber perpendicular to the molecular beam. A vacuum window is located right above the exit of the decelerator in both detection chambers and allows for placement of a PMT. For REMPI, the MCP is placed along the molecular beam axis, 20-30 cm downstream of the detection region. Either a simple two plate TOFMS is installed immediately after the decelerator or an electrostatic trap can be turned to a TOFMS voltage configuration to accelerate the ionized molecules. For all experiments, we use a 40 mm two-stage

chevron configuration MCP. The molecule signal is detected as a current, which is then converted to a voltage using a transimpedance amplifier and the resulting voltage versus time signal for the desired mass channel is integrated.

The laser system used for detection is a pulsed dye laser (Sirah PrecisionScan PRSC-D-24SL with a doubling stage) pumped by 532 nm light from a doubled-pulsed Nd:YAG laser (QuantaRay Lab170). Conveniently, we detect ND₃ using a 2+1 REMPI process (see section 3.1.2) and only one laser wavelength is required. To detect OH, on the other hand, is slightly more complicated since it uses a 1+1' system. The first photon at 282 nm is generated using the dye laser. The second photon, 118 nm, is generated by frequency tripling 355 nm light from a tripled Nd:YAG in a xenon-argon gas cell. This procedure covered in Chapter 6. The two beams are counter propagated through the chamber as shown in Figure 2.15.

2.5.3.1 Pin Decelerator

The pulsed-pin Stark decelerator in our lab has 150 pairs of pins. Typically ± 12 kV is applied to the pin in this decelerator. In PPSD, in order to create the correct shape potential well, the first pair of pins must start grounded. This creates a defined potential well between the first two sets of pins. A defined well does not exist before the first pair of pins. This means a decelerator with 150 pin pairs has 149 deceleration stages. Figure 2.18 shows a CAD drawing of the decelerator from two angles and has pin-to-pin spacing labeled in mm. Each pin has a diameter of 0.125 inches (3.175 mm) and has a center-to-center spacing of 2.04 inches (5.1816 mm) with its corresponding pin in the transverse direction. The center-to-center spacing in the longitudinal direction between adjacent (they will be perpendicular to each other) pins is 0.215 inches (5.461 mm).

The pins are set into four mounting rods (see Figure 2.1a), and each rod is connected to a Behlke HTS-201-03-GSM switch. Four digital logic signals are required to control the four switches. Although paired rods switch at the same time, the polarity of the Behlke switches means inverted control signals are required for the negative electric potential rod of each pair. Thus four control signals and four switches are required to run a pulsed decelerator.




Figure 2.18: A CAD drawing of the pulse-pin decelerator along the molecular beam axis (left) and perpendicular to the molecular beam axis (right). All measurements are in mm. Taken from Figure 4.9 in [155]




Figure 2.19: The four electrode electrostatic trap used with the pulse-pin Stark decelerator. The typical voltages used on each electrode are given and this creates a quasi-quadrupole trap potential with a larger gradient in the x and y direction than in the z direction.

At the end of our pulsed-pin Stark decelerator, we have a four electrode trap (Figure 2.19). It is similar to the quadrupole trap demonstrated in [85], but with an additional 2 mm slit cut into the last electrode. This slit was designed to assist with extracting ionized molecules when the detection lasers are scanned in the vertical (y) direction. Typically, the trap is run with the voltages [-10, 10, 10, -10] kV on the four electrodes starting with the upstream electrode. This creates a quasi-

quadrupole trap. The electrostatic trap that most closely resembles a pure quadrupole [84, 114] does not have optical access. By having two center electrodes, instead of just one, an ionization laser can enter perpendicular to the molecular beam axis. The flexibility of the four electrodes also allows for a loading configuration, referred to as “half on,” which is achieved using [0, 10, 0, -10] kV on the electrodes. This produces a Stark potential hill with a minimum offset upstream from the trap center, which is used as an additional deceleration stage before the trap is turned on (Figure 2.20). Previously, the two center electrodes were spaced 5 mm apart so that rubidium atoms could be loaded from the side between the two electrodes as part of a ND_3+Rb cotrapped experiment [117]. Modifications have been made so that the spacing is only 3 mm. This increases the trap depth. Additionally, in previous work, the trap was operated at ± 8 kV, it is currently operated at ± 10 kV. Lastly, a 5 mm hole was cut in the last trap electrode so that the Rb atoms can still be loaded into the trap, but now from the back of the trap instead of from the side [155].

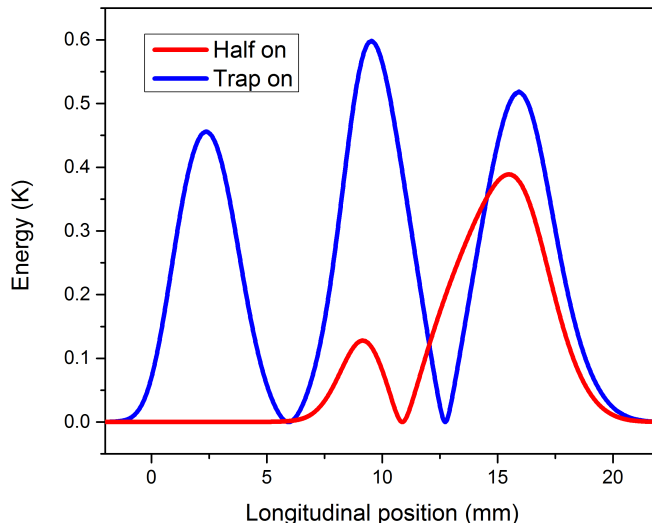


Figure 2.20: The longitudinal on axis potentials for the trap on (blue) and half on (red) configurations of the trap. The half on stage is used as an extra deceleration stage during trap loading. The center of the trap is located near 13 mm and the extra potential energy hill in the trap on configuration is created by the grounded decelerator electrodes and the first electrode of the trap.

2.5.3.2 Ring Decelerator

The ring decelerator in our lab has 624 rings. Each ring is made of 0.04 inch (1.016 mm) diameter tantalum wire. The rings are each shaped like tennis rackets (Figure 2.8a) and held in the rods with two set screws. The decelerator has a periodicity of eight, with 78 rings per rod. In order for the decelerator to function properly and hold the correct voltages, the surface quality and cleanliness of the rings is very important. The rings are mechanically polished on a buffing wheel, electropolished, and then thoroughly cleaned before they are instilled into the decelerator. Figure 2.21 shows various steps of assembling the decelerator. Each decelerator rod is made of many sections each with 10-20 rings. Once the rings are attached to each section, the rods are assembled and the eight rods placed into their mounts. The decelerator is placed in a mounting cage, which keeps it ridged and assists with alignment and installation inside the UHV chamber. Once assembled in the cage, the rings must be aligned to ensure all rings are centered on the decelerator axis and are parallel to each other. Careful alignment is completed with Teflon coated tools so that the rings are not scratched in the alignment process.

After installation into the vacuum chamber, the full chamber with the decelerator inside was baked to reduce the background pressure in the chamber. The decelerator was voltage conditioned using 200 μ s pulses up to ± 10.5 kV. However, we believe better conditioning resulted from using the amplifiers on the decelerator and slowly increasing the voltage of chirped sine-wave over the course of several weeks.

An eight ring trap has also been designed for our decelerator (Figure 2.22). It maintains the same size rings and ring spacing as the decelerator, but each ring is electrically isolated. It is mounted via a stainless steel plate and Macor standoffs. This allows for more flexible voltage waveforms in the trapping region, since the entire decelerator does not need to remain at high voltage while the trapping potentials are on. The trap can also be used in a TOFMS configuration with the addition of a voltage divider box that evenly divides down the voltage from 1 kV to ground in eight steps. Although this trap provides significant flexibility for trapping studies, it also requires





Figure 2.21: The steps involved in assembling a ring decelerator. A) Rings are attached to rod sections. B) Each rod is comprised of 4 sections. C) The sections are assembled into rods. D) The rods are arranged in the correct order such that the ‘legs’ of each ring form a helical pattern. E) Align the rings so that they are all concentric about the same axis. F) Fully assembled decelerator in its support rack.

eight amplifiers in addition to the eight already on the decelerator, making it a significant challenge to have sixteen analog high-voltage amplifiers operating at the same time. With the eight electrically isolated rings attached to the end of our decelerator, we were able to use REMPI detection between the fourth and fifth rings of the trap as a TOFMS and detect decelerated molecule signal. However, the full trapping configuration still remains to be realized. For the studies discussed in Chapter 4, the trap was removed from the experiment in favor of a simple TOFMS configuration made of two parallel plates with a circular 5 mm diameter center opening. Typically 1 kV is applied to the upstream plate and the second plate is grounded.



(a) CAD drawing of ring trap



(b) Ring trap

Figure 2.22: The ring trap

2.6 Decelerator comparison

Choosing the appropriate decelerator geometry and running mode for a given application will depend on the molecule in question and initial and final velocities. The many details of Stark deceleration will be discussed in subsequent chapters, but here we present a general overview and a comparison of the two decelerators constructed in our lab.

The molecule Stark energy and the decelerator geometry and voltage will set the gradient of the Stark potential well experience by the molecules. The initial and final velocities of the beam and length of the decelerator will determine the linear acceleration required (Equation 2.27). The initial velocity of the beam will be set by the source and the final velocity is often set by the limits of the experiment. The biggest trade off that must be decided on is length versus voltage. A longer decelerator means lower voltages and accelerations can be used, but it increases the costs and difficulties to build and maintain. Sometimes, this is unavoidable due to the molecule being decelerated, for heavy molecules that do not have linear Stark shifts, a longer decelerator may be the only solution. Generally, a longer decelerator of any kind will have an increased capacitance and will require sourcing more current to keep the decelerator at the right voltages. A longer pin decelerator may lead to more transverse losses unless a mode such as S=3 is used. Additionally, for any decelerator geometry, the more electrodes there are, the more difficult it is to build, align, and voltage condition. Higher voltages, on the other hand, increases the potential well gradient and the accelerations it can accommodate, but increase the difficulty of the electronic and voltage conditioning. How much voltage can be supported across two electrodes in vacuum is set by their separation distance and surface quality.




Figure 2.23: On-axis longitudinal potentials for PPSD (green), PRSD (red), and TWSD (blue) modes in our pin and ring decelerators. For both modes of the ring decelerator, both ± 7 kV (solid) and ± 12 kV (dashed) potentials are shown. All potentials are for ND_3 .

Table 2.2: Properties of the two decelerators used.

Geometry	Crossed pins	Rings
Number of electrodes	150	624
Stage length	5.461 mm	2.032 mm
Total length	0.81 m	1.27 m
HV Electronics	Switches	Switches or Analog amplifiers

It is difficult to universally compare Stark deceleration methods since the efficiency is highly dependent on the exact geometry and voltages applied. Each decelerator is designed and built to fit the space and application it is to be used for. Just for reference, here, we compare the two physical geometries and three modes used in this dissertation, but these comparisons do not carry over into other pin/ring decelerators. Table 2.2 gives a summary of the two physical decelerators and Figure 2.23 shows the on-axis longitudinal potentials for PPSD (green), PRSD (red), and TWSD (blue) modes in our decelerators. For both modes of the ring decelerator, both ± 7 kV (solid) and ± 12 kV (dashed) potentials are shown. ± 7 kV is used in this dissertation, but ± 12 kV is also shown so that the wells for the same maximum voltage can be compared. It is straightforward to see the effect of geometry and voltage from these plots. The shape of the bottom of the potential well is very different and longitudinally longer from the pin decelerator to the ring decelerator. Even for the same maximum voltages, the well for the pin decelerator is much deeper. However it has significantly fewer stages and more loss mechanisms as mentioned above. Between the two ring decelerator modes, PRSD has a smaller Stark potential well, but with a larger gradient, which plays an important role in its efficiency. A well with steeper gradients can experience more acceleration before it can no longer successfully decelerate any molecules.

Chapter 3

Molecules of Interest

In this dissertation, we decelerate deuterated ammonia (ND_3) and hydroxyl (OH) in weak-field-seeking states. This chapter presents a review of some basic molecular structure important to understanding the Stark deceleration and detection of these molecules. We then discuss the structure of ammonia and OH and how they behave in electric fields.

Molecules can be excited electronically, vibrationally, or rotationally. This provides more degrees of freedom and a significantly more complex energy-level structure than in atoms. As molecules vibrate, this changes the internuclear spacing between the constituent atoms, which affects the moments of inertia of the molecule and its rotational energies. Similarly, centrifugal forces during rotation distort the equilibrium distances between atoms, which affects its vibrations. This coupling between rotational and vibrational energy leads to rich rovibrational energy level structures within molecules. While Stark decelerated molecules remain in the rovibrational ground state, the excited state vibrational and rotational energy levels are important for state-selective detection of these molecules.

The Schrödinger equation for even the simplest molecule, two protons and an electron, cannot be solved analytically. To allow for analytical solutions, the Born-Oppenheimer approximation is made, which allows the electronic and nuclear motions to be decoupled and separated. This approximation is based on the relative speeds of the nuclei compared to the electrons. Since the masses of the nuclei are much greater than the masses of the electrons, the approximation assumes that the electrons instantly adapt to changes in the position of the nuclei. The molecular wave

function, Ψ_{mol} , can then be separated into electronic and nuclear components

$$|\Psi_{mol}\rangle = |\Psi_{elec}\rangle |\Psi_{nuc}\rangle. \quad (3.1)$$

This solution has a total energy given by $E_{mol} = E_{elec} + E_{nuc}$.

The Hamiltonian for the molecule can now be considered to have separate electronic and nuclear components. An additional relativistic component is not considered here, but can be treated perturbatively. The electronic wave functions can be solved for different nuclear configurations. The nuclear component can be additionally divided into motion about the equilibrium position (vibrations) and motion about the center of mass (rotations). This results in the Hamiltonian

$$\begin{aligned} H_{mol} &= H_{elec} + H_{nuc}, \\ H_{mol} &= H_{elec} + H_{vib} + H_{rot}, \end{aligned} \quad (3.2)$$

where the energy of the molecule is given by $E_{mol} = E_{elec} + E_{vib} + E_{rot}$. For the most part, these energies occur on different scales. The electronic energy levels are the largest and have energy separations that occur in the visible or ultraviolet region of the electromagnetic spectrum. The vibrational energy-level spacings are typically found in the infrared region, and rotational energy-level spacing are in the microwave region.

Electronic states of molecules have spectroscopic term symbols just like those of atoms. They are dependent on the symmetry of the molecule and how the many sources of angular momentum couple together. Only after analyzing the many angular momentum terms can the term symbols and selection rules be determined. For diatomic molecules, molecular term symbols are written as $^{2S+1}\Lambda_{\Omega}$, where $2S + 1$ is the spin multiplicity. Λ is the projection of the total orbital angular momentum on the internuclear axis. Λ takes on the values of 0, 1, 2... which correspond with the upper case Greek letters Σ , Π , and Δ respectively. Typically, the lowest electronic state is labeled with the letter X and subsequent energy levels are labeled using A, B, C,..., with A being the first excited state above the ground state. If the molecule has states with different spin multiplicity, those states are labeled using lower case letters (a, b, c,...) in order to differentiate them from those that can be accessed from the ground state by an electric dipole transition.

For the Stark deceleration experiment discussed here only the ground rovibrational states are significantly populated. Here, we present a quick discussion of rotational states in molecules, which is primarily based on [156], and the rotational states of ammonia and OH in the ground state, which will aid in selecting a good detection scheme for these molecules.

3.0.1 Rotational states

To first order, we consider the molecule to be rigid. The moment of inertia of a rigid system about an axis q is

$$I_q = \sum_i m_i (r_i(q))^2, \quad (3.3)$$

where m_i is the mass of each atom and $r_i(q)$ is the distance of the atom to the axis q . Since we are considering a rigid system, the classical kinetic energy of a rotating body can be used, and the energy of the system is

$$E = \frac{1}{2} \sum_q I_q \omega_q^2 = \frac{1}{2} \sum_q \frac{J_q^2}{I_q}, \quad (3.4)$$

where ω_q is the angular frequency and J_q is the angular momentum about the axis q . This leads to the Hamiltonian for rotational energy in Cartesian coordinates to be

$$H = \frac{J_x^2}{2I_x} + \frac{J_y^2}{2I_y} + \frac{J_z^2}{2I_z}. \quad (3.5)$$

In order to determine good quantum numbers for the rotational states of a molecule, Equation 3.5 can be simplified based on the symmetries of the molecule of interest. Ammonia is a symmetric top molecule so two of the moments of inertia are identical, and Equation 3.5 becomes

$$\begin{aligned} H &= \frac{J_x^2 + J_y^2}{2I_{\perp}} + \frac{J_z^2}{2I_{\parallel}} \\ &= \frac{J^2}{2I_{\perp}} + \left(\frac{1}{2I_{\parallel}} - \frac{1}{2I_{\perp}} \right) J_z^2, \end{aligned} \quad (3.6)$$

where the total angular momentum has been re-expressed using $J^2 = J_x^2 + J_y^2 + J_z^2$. We now recall that the eigenvalues of the J^2 operator are $J(J+1)\hbar$ and those of J_z are $m\hbar$; however, in the case of molecules, the quantum number K is used instead of m , since M is reserved for the projection of J onto the laboratory frame. The good quantum numbers are now:

- $J = 0, 1, 2, \dots$
- $K = -J, -J + 1, \dots, J - 1, J$
- $M = -J, -J + 1, \dots, J - 1, J$.

Due to the cylindrical symmetry of a symmetric rotor, the molecule's rotational energy depends on only the magnitude of K and thus K may be specified as $K = 0, 1, \dots, J$. This comes about because of the J_z^2 term in Equation 3.6.

Rearranging Equation 3.6 gives

$$\begin{aligned}
 H &= \frac{1}{2I_{\perp}} (J^2 - J_z^2) + \frac{J_z^2}{2I_{\parallel}} \\
 &= B_0 (J^2 - J_z^2) + A_0 J_z^2 \\
 &= B_0 J^2 + (A_0 - B_0) J_z^2,
 \end{aligned} \tag{3.7}$$

where the constants A_0 and B_0 are rotational constants specified by

$$A_0 = \frac{1}{2I_{\parallel}} \text{ and } B_0 = \frac{1}{2I_{\perp}}. \tag{3.8}$$

The energy of the rotational states are then specified by

$$E(J, K) = B_0 J(J + 1) + (A_0 - B_0) K^2. \tag{3.9}$$

Here, M has not played a role, but it will once an electric field is applied.

OH, on the other hand, is a linear rotor. In the case of a linear rotor the angular momentum vector is perpendicular to the internuclear axis, which results in $A_0 = 0$ and $K=0$. The rotational energies for the field free linear rotor are specified by

$$E(J) = B_0 J(J + 1). \tag{3.10}$$

With this understanding of rotational structure, we now examine the rotational levels of the two molecules of interest and how the ground rovibrational energy behaves in the presence of electric fields.

Table 3.1: Properties of ND₃

N-D bond length	1.02 Å
Pyramid equilibrium height	0.38 Å
Umbrella angle	68°
Electric dipole moment (μ_e)	1.48 Debye [161]

3.1 Ammonia

Ammonia has been widely studied because of its symmetry and importance in agriculture and astronomy [157–159]. In particular, it is a useful tool for examining temperatures and densities in the interstellar medium [160]. Ammonia is often used in Stark deceleration experiments due to its ease of implementation, large electric dipole moment, low mass, and stability at room temperature and pressure. This is unlike OH and NH, radicals that are also used in decelerator experiments, which must be created *in situ* using discharge plate. Additionally, ammonia can be quantum state selectively detected using only one laser wavelength, which makes detection straightforward to implement. Both deuterated ammonia (ND₃) and regular ammonia (NH₃) are used for Stark deceleration. However, ND₃ is more commonly used because of its smaller inversion splitting, which leads to a larger Stark shift at a given electric field. Here, when we refer to ammonia, we will be discussing ND₃ unless otherwise noted.

Ammonia is a polyatomic symmetric top molecule. It has trigonal pyramidal structure with the nitrogen at the peak of the pyramid and three hydrogens that form the base. Table 3.1 gives some basis properties of deuterated ammonia. An inversion splitting arises from the nitrogen atom tunneling through the plane of the three hydrogen atoms. This gives Ammonia a double well potential energy surface. Although inversion is a type of vibration, it occurs in the microwave regime at a scale even smaller than those of rotations.

An additional symmetry results from the symmetric or asymmetric exchange of the three identical nuclei. In the *ortho* state, K is a multiple of three, while in the *para* state, K is not a multiple of three. Conversion between the *ortho* and *para* states is not common since the nuclear spin does not couple to the molecular rotation.

Figure 3.1 shows the rotational energy levels of the electronic and vibrational ground state with the J , K and *ortho/para* structure labeled. The ground state of ammonia has the term symbol $\tilde{X}^1A'_1$.




Figure 3.1: The rotational levels of ND_3 in the electronic and vibrational ground state. The J value is labeled on top of each state and they are sorted by K , which dictates *ortho* and *para* ladders. The inversion splitting on each rotational level is not shown. This figure is taken from Figure 2.6 in [95].

3.1.1 Stark effect in ND_3

With no external electric fields, the ground state ND_3 does not have a permanent electric dipole moment. As the molecule undergoes inversion, the effective dipole moment cancels out. It is only when an external electric field is applied that an electric dipole moment is induced. Since ammonia is a symmetric top molecule, its Stark shift, U_{stark} , in a DC electric field is given by [162]

$$U_{stark} = \mp \left(\frac{\Delta U_{inv}}{2} \right) \pm \sqrt{\left(\frac{\Delta U_{inv}}{2} \right)^2 + \mu E \frac{MK}{J(J+1)}}, \quad (3.11)$$

where ΔU_{inv} is the inversion splitting, μ is the electric dipole moment, and E is the magnitude of the electric field. J , K , and M note the rotational state. Now that an external field has been applied, M , the projection of J on to the electric field, plays a role. The signs in Equation 3.11 correspond to the dipole moment being anti-aligned (top) or aligned (bottom) with the electric field. This results in the Stark energy increasing (decreasing) with the magnitude of the electric field, which leads to weak-field-seeking (strong-field-seeking) states. The Stark shift is quadratic

at weak electric field strengths, but becomes linear at high electric field strengths. Figure 3.2 shows the Stark shift for both ND_3 and NH_3 in the $J=1, K=1$ ground state. Here the affect of the inversion splitting between ND_3 and NH_3 becomes apparent. The inversion splitting, ΔU_{inv} , is 0.79 cm^{-1} for NH_3 and 0.053 cm^{-1} for ND_3 . The smaller inversion splitting for ND_3 leads to a larger Stark shift, and explains why it is used more often for Stark deceleration experiments. More energy can be removed from ND_3 than NH_3 for the same electric field and thus lower electric fields can be used overall.




Figure 3.2: The Stark energy as a function of electric field for ND_3 and NH_3 in the $J=1, K=1$ state. This figure is adapted from Figure 2.7 in [95].

3.1.2 ND_3 REMPI

The ionization potential of ND_3 is 10.07 eV, so it can be ionized with one VUV photon, but that process would not be state selective. In laboratory experiments, state-selective detection of ND_3 is commonly done using a 2+1 REMPI scheme, with one laser at 317 nm [163–166]. A 3+1 REMPI scheme for ammonia has also been demonstrated [167], but is not as commonly used. The 2+1 scheme is convenient because it requires only one laser wavelength, which can be produced by a doubled dye laser, at 317 nm. Figure 3.3 depicts the states used in this detection scheme. Ammonia starting in the upper inversion doublet (the one that is weak-field-seeking and can be Stark decelerated) is excited to the intermediate \tilde{B}^1E'' , $v_2 = 5$ Rydberg state state via two 317 nm

photons. There are several allowed ionization pathways from different rotational levels that can be used in the $\tilde{B}^1E'' \leftarrow \tilde{X}^1A_1'$ transition [168]. The \tilde{B}^1E'' state is planar and thus does not have an inversion splitting. From here, ND_3 can then be ionized with an additional 317 nm photon.




Figure 3.3: The 2+1 REMPI scheme for detecting ND_3 .

3.2 OH

OH is an important radical in atmospheric and environmentally relevant reactions [169,170]. It is also an interesting molecule in crossed electric and magnetic fields since it has both an electric dipole moment (1.67 Debye) and a magnetic dipole moment ($2\mu_b$, where μ_b is the Bohr magneton). Since OH is a diatomic molecule, before we discuss the OH structure, it is useful to review Hund's coupling cases, which categorizations of how the idealized electronic and orbital angular momenta of diatomic molecules couple together

3.2.1 Hund's Coupling Cases

Determining the Hund's case for a given electronic state will help guide what quantum numbers are best suited for the given state of the molecule. There are typically five cases (a)-(e), but

even then these do not describe all situations. For example, OH in strong electric and magnetic fields is describe by a “Hund’s case (x)” [171]. Here we will discuss cases (a) and (b), which are the two most common cases.

There are four sources of angular momentum in a molecules: the spin of the electrons (\mathbf{S}), the electronic orbital angular momentum (\mathbf{L}), the rotation of the nuclei (\mathbf{R}), and the nuclear spin (\mathbf{I}). The total angular momentum, not including \mathbf{I} , is \mathbf{J} . If the nuclear spin is included, the total angular momentum is \mathbf{F} . However, nuclear hyperfine effects resulting from the coupling of the nuclear spin with the magnetic field from another angular momentum component is very small and not considered here. A summary of the angular momentum vectors is:

- \mathbf{L} is the electronic orbital angular momentum
- $\mathbf{\Lambda}$ is the component of \mathbf{L} along the internuclear axis
- \mathbf{S} is the electronic spin angular momentum
- $\mathbf{\Sigma}$ is the component of \mathbf{S} along the internuclear axis
- \mathbf{J} is the total angular momentum
- $\mathbf{\Omega}$ is the vector sum of \mathbf{L} and \mathbf{S} along the internuclear axis
- \mathbf{K} is the vector sum of $\mathbf{\Lambda}$ and \mathbf{N}
- \mathbf{N} is the total angular momentum but not including the electron spin, $\mathbf{N} = \mathbf{J} - \mathbf{S}$
- \mathbf{R} is the rotational angular momentum of the nuclei, $\mathbf{R} = \mathbf{N} - \mathbf{L}$.

In Hund’s case (a), the electronic state has large spin-orbit coupling and a strong electrostatic coupling of \mathbf{L} to the internuclear axis. As a result of the spin-orbit coupling, \mathbf{S} is also strongly coupled to the internuclear axis. This makes $\mathbf{\Omega} = \mathbf{\Lambda} + \mathbf{\Sigma}$ a good quantum numbers for this case. The good quantum number for this case are $\mathbf{\Omega}$, $\mathbf{\Sigma}$, $\mathbf{\Omega}$, and \mathbf{J} . $^2\Pi$ is a typical Hund’s case (a) state.

In Hund’s case (b), the spin-orbit coupling is weak. The orbital angular momentum still couples to the internuclear axis but the spin does not. Instead, \mathbf{R} , the rotational angular momentum

of the nuclei, couples to the internuclear axis giving $\mathbf{R} + \mathbf{L} = \mathbf{N}$. In this case, the total angular momentum \mathbf{J} is given by the sum of $\mathbf{N} + \mathbf{S}$. $^2\Sigma$ states and $^2\Pi$ states with large \mathbf{J} are typical Hund's case (b) states.

3.2.2 OH structure

OH is a low mass, open shell radical with a strong electric and magnetic dipole moments. It has a large rotational constant, which results from the small moment of inertia. The $^2\Pi_{3/2}$ ground state has non-negligible mixing with the spin-orbit split $^2\Pi_{1/2}$ state. A small (~ 1.7 GHz) Λ -doublet splitting that arises from the coupling of the angular momentum projection Λ and nuclear rotation of excited electronic states. Similar to the case in ammonia, the vibrational levels of the electronic ground state of OH are too large to be significantly populated at the temperatures in the supersonic expansion use in these experiments. Thus, we are concerned primarily with rotational splittings and spin-orbit coupling within the ground state.

The molecular orbitals for ground state OH is

$$(1s\sigma)^2(2s\sigma)^2(2p\sigma)^2(2p\pi)^3. \quad (3.12)$$

OH is an open shell radical since its highest occupied molecular orbital is not completely filled. The spin and orbital angular momentum of this unpaired electron will dictate the properties of the electronic ground state of the molecule. In the ground state, the valence electron is in a π orbital, so $\Lambda = \pm 1$ and $\Sigma = \pm 1/2$, and the electronic ground state of the OH molecule is $^2\Pi_{3/2}$. The $^2\Pi$ electronic ground state of OH is primarily Hund's case (a). In the first excited state, a $2p\sigma$ electron is promoted to the $2p\pi$ state giving the electron configuration of

$$(1s\sigma)^2(2s\sigma)^2(2p\sigma)^1(2p\pi)^4. \quad (3.13)$$

The lone σ electron has $\lambda = 0$ and no projection on the internuclear axis. This results in the first electronic excited state having Σ character and is best described by Hund's case (b).

Good quantum numbers for the ground state Hund's case (a) state are Λ , Σ , $\bar{\Omega}$, J , M_j , and ϵ . $\bar{\Omega}$ is the vector sum of L and S and thus can have the values $\pm\frac{3}{2}$ and $\pm\frac{1}{2}$. The total angular

momentum is $J=3/2$. $\bar{\Omega} = |\Omega|$ which gives the two spin-orbit split manifolds of ${}^2\Pi_{3/2}$ and ${}^2\Pi_{1/2}$. ϵ is the symmetry of the state and can either be positive or negative, it is typically denoted using either e ($\epsilon = +1$) or f ($\epsilon = -1$). λ -doubling breaks the degeneracy of the rotation about the internuclear axis. Lastly, there is one more quantum number \mathcal{P} , the parity of each state. \mathcal{P} arises because diatomic molecules possess reflection symmetry. This leads to $\mathcal{P} = \pm 1$ as good quantum numbers where $\mathcal{P} = +1$ is for an even function under spatial inversion and $\mathcal{P} = -1$ is for an odd function under spatial inversion.

For the ${}^2\Pi_{3/2}$ rovibration ground state the λ -doublet splitting is $\Delta \approx 1.667$ GHz (0.0556 cm^{-1}). The rovibration ground state of OH is $J = \frac{3}{2}$, $\epsilon = e$ and the $J = \frac{3}{2}$, $\epsilon = f$ state lies Δ above it. Figure 3.4 shows the rotational structure in the electronic and vibrational ground state of OH.

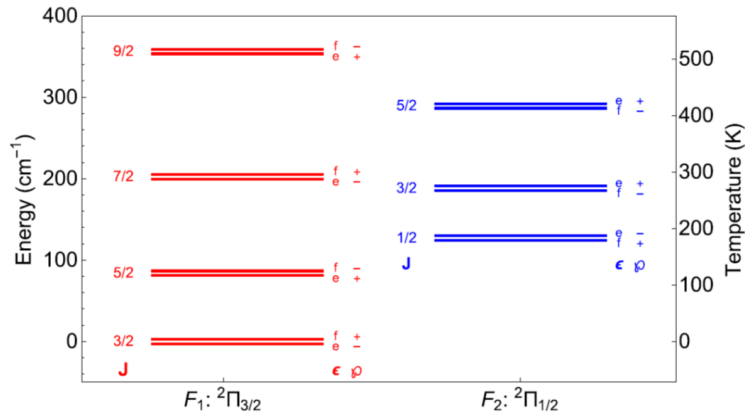


Figure 3.4: The rotational structure in the electronic and vibrational ground state of OH. The λ -doublet splitting has been exaggerated for clarity. \mathcal{P} is the parity of each states.

3.2.3 OH in electric fields

The Stark shift in OH (similar to Equation 3.11) is

$$U_{stark} = \mp \left(\frac{\Delta}{2} \right) \pm \sqrt{\left(\frac{\Delta}{2} \right)^2 + \mu E \frac{M\Omega_{eff}}{J(J+1)}}. \quad (3.14)$$

A detailed derivation starting from the Hamiltonian for the Stark energy in the fixed lab frame and including a transformation between the fields in the laboratory frame and the dipole moment

in the molecule frame can be found in [155]. A plot of the stark shift for the $J = 3/2$ rovibrational ground state is shown in Figure 3.5. A further discussion of OH in crossed electric and magnetic fields and 1+1' REMPI detection of OH can be found in Chapter 6.




Figure 3.5: The Stark splitting in OH.

Chapter 4

Pulsed Ring Stark Deceleration

Pulsed ring Stark deceleration combines the cylindrical symmetry of TWSD and the relative electronic simplicity of PPSD. Its discrete voltages make it easier to implement than the analog high-voltage amplifiers required for TWSD. Although TWSD is the “premier” Stark deceleration technique because of its continuous nature, in certain situations PRSD is more efficient than TWSD due to an increased longitudinal potential gradient. This is particularly important when there is a limit on the maximum voltage that can be used on the electrodes or on the length of the decelerator.

In this chapter, we present an experimental realization of a ring-geometry Stark decelerator using both continuously varying (TWSD) and discrete voltages (PRSD). Using the pulsed-ring method, we demonstrate a newly accessible low velocity regime for moderate peak voltages on the ring electrodes. A comparison of experimental and simulated results between traveling-wave and pulsed-ring Stark deceleration is presented along with a simple model for determining when each mode is more efficient. Section 4.1 of this chapter is adapted from [172] and some material maybe a review of what has already been discussed. The following sections provide a more in depth analysis of how to implement PRSD and the transverse dynamics of the Stark potential well.

4.1 Pulsed-Ring versus traveling-wave Stark deceleration

4.1.1 Introduction

Stark deceleration uses spatially inhomogeneous and time-varying electric fields to decelerate neutral polar molecules, and produces a beam of slow molecules with a tunable average final

velocity. These beams can be used for high-resolution spectroscopy [99], collision physics [100], measuring radiative lifetimes of molecules [101,102], and tests of fundamental physics [10,104,105]. The molecules from a Stark decelerator can also be loaded into electrostatic [84–86], magnetic [87,88], and ac traps [89–91], or used in crossed beam [75–78], beam-trap [92,93], and co-trapped experiments [94]. Such experiments are often limited by the number of successfully decelerated molecules. Improvements to the decelerator efficiency, the percentage of molecules in the initial pulse that are successfully decelerated, would benefit these types of studies.

The first Stark decelerator, referred to in this paper as the pulsed-pin Stark decelerator (PPSD), was constructed in 1999 by Bethlem *et al.* [107]. It used high-voltage switches to alternate between two static voltage configurations on pairs of crossed pin electrodes. Switching between these voltage configurations produces a discretely moving Stark potential well to decelerate molecules. Molecules only within a range of positions and velocities will be decelerated. This portion of phase space is called the phase-space acceptance. Ideally, it would be well filled with molecules. PPSD has well-characterized instabilities and loss mechanisms that are enhanced at low final velocities [139]. Particularly, coupling between the longitudinal and transverse motion leads to unstable regions within the phase-space acceptance. The decrease in phase-space acceptance drastically reduces the efficiency when decelerating to low velocities required for trapping, or when spending long times at velocities <100 m/s in the decelerator. Although different timing schemes and operating modes can improve the performance of the traditional Stark decelerator [117–119], the crossed-pin electrode geometry fundamentally limits the efficiency due to longitudinal and transverse coupling.

The loss present in pin decelerators can be mitigated by using ring-shaped electrodes and a fundamentally different mode of applying voltages to the electrodes: chirped sinusoidally varying voltages. This produces a traveling-wave potential well that continuously guides or decelerates the molecular beam, as opposed the pulsed nature of PPSD. Traveling-wave Stark deceleration (TWSD) has two characteristic differences compared to PPSD: the cylindrically symmetric electrode geometry and continuously, as opposed to discretely, varying voltages, which produce a genuine moving potential well. This leads to an inherently stable deceleration process that has

a true three-dimensional (3D) trapping potential and a well-filled phase-space acceptance. These factors allow for more efficient deceleration due to the continuous and cylindrically symmetric nature of this deceleration mode. Meek *et al.* [132] first demonstrated the idea for spatially modulated potential wells on a microstructured chip used to guide and then decelerate metastable CO molecules [133,134]. The first use of macroscopic traveling-wave Stark potentials with a ring-geometry decelerator was experimentally shown in 2010 by Osterwalder *et al.* [121] using a beam of CO, which was decelerated from 288 m/s to 144 m/s.

While TWSD has many advantages, it is more challenging to implement than PPSD because it requires high-voltage analog amplifiers with demanding specifications to generate the moving Stark potential wells. The sinusoidal voltage must have an initial frequency such that the Stark potential well containing the molecules moves with the same initial velocity as the molecular beam. The frequency of the sinusoidal voltages then must be smoothly chirped down to a frequency such that the velocity of the Stark potential well formed by the voltages matches the desired final beam velocity. Thus, the amplifiers require a large bandwidth to decelerate a typical molecular beam to rest. The chirp rate of the sinusoidal voltage will be determined by the initial and final velocities and the length of the decelerator. For a fixed ring geometry, the magnitude of the maximum allowable deceleration is set by the bandwidth of the high-voltage amplifiers, maximum output voltage of the high-voltage amplifiers, and length of the decelerator.

Early TWSD experiments were limited by the lack of availability of high-voltage amplifiers with a large bandwidth. The amplifiers used to generate the voltages for the electrodes in the first TWSD experiment operated at ± 8 -10 kV with 10-30 kHz sinusoidal voltages [121,122]. The limited bandwidth of the amplifiers meant that the decelerator could achieve only moderate deceleration and the molecules could not be brought to rest. Commercial amplifiers (Trek 5/80), which are capable of a larger frequency range (0-60 kHz) and outputs up to ± 5 kV, make it possible to decelerate a supersonic beam down to rest with a decelerator that is several meters long. For example, a beam of SrF seeded in Xe with an initial velocity of 290 m/s would require a ~ 5 m long decelerator [123]. Amplifiers with slightly limited maximum voltages, such as these, are

particularly useful for TWSD of heavy molecules used in precision measurements, since the electric fields produced are low enough to keep the molecules in the weak-field seeking regime, where the Stark energy increases with increasing electric field [124]. To date, YbF [125] and SrF [123] have been decelerated using TWSD and the longest demonstrated decelerator (4 m) slowed SrF from 290 to 120 m/s using ± 5 kV amplifiers [126].

To overcome the limitations of PPSD and the demanding amplifier specification of TWSD, a combination of both Stark deceleration methods can be used. In this combination technique, a pulsed-pin Stark decelerator decelerates molecules to longitudinal velocities around 100 m/s, and then a traveling-wave Stark decelerator decelerates the molecules to rest [127]. This scheme is able to take advantage of the cylindrical symmetry of the ring geometry at low velocities to avoid the low velocity losses of the pulsed-pin Stark decelerator, but it relaxes the requirement to have either extremely high maximum voltages or a very long decelerator to reach electrostatically trappable conditions. Using this combination technique, NH_3/ND_3 and CH_3F have been trapped, manipulated, and adiabatically cooled within the ring decelerator [127–130].

Complete TWSD of a supersonic beam moving at 300–450 m/s to rest has been challenging to achieve due to electronic and hardware requirements. An alternative running mode for operating a ring-geometry decelerator has been proposed by Hou *et al.* [141], and uses pulsed voltages on ring electrodes. The duration of the pulses is varied in an identical manner to that of PPSD. This pulsed-ring Stark deceleration (PRSD) method takes advantage of the cylindrical symmetry provided by the rings and is simpler to implement than TWSD, since it uses commercial high-voltage switches instead of high-voltage analog amplifiers. In this paper, we present experimental results of this new PRSD scheme in comparison with TWSD, examine conditions where it is advantageous to run PRSD over TWSD, and discuss a simple method for determining which mode is preferable based on the desired final molecule packet properties.

4.1.2 Experiment

The experimental setup for demonstrating both deceleration modes is shown in Fig. 4.1. Both experiments begin with a piezoelectric-actuated pulsed valve that creates a pulsed supersonic beam of 2% fully deuterated ammonia (ND_3) in krypton. A skimmer ~ 20 cm downstream of the valve collimates the molecular beam and allows for differential pumping between the source and decelerator chambers. The decelerator begins 3 mm behind the skimmer, and consists of 624 rings mounted in eight longitudinal stainless steel rods such that every eighth ring is electrically connected (similar to [121]). Each ring is made of 1.02 mm diameter tantalum wire, and has a 4 mm inner diameter and center-to-center spacing of l , which for the decelerator discussed here is $l = 2.03$ mm. The full decelerator is just over 1.25 meters long. For the work presented here, ND_3 molecules are decelerated from 415 m/s down to final velocities of 220 m/s for TWSD and 150 m/s for PRSD. ND_3 molecules at the end of the decelerator are ionized in a 2+1 REMPI scheme using 317 nm photons [166]. Two stainless steel plates, which are electrically isolated from the decelerator, are mounted after the rings. They form the time-of-flight mass spectrometer and accelerate ionized molecules into the microchannel plate detector, where the measured current is proportional to the number of ionized molecules.




Figure 4.1: A schematic of the experimental setup which includes, from left to right, a pulsed valve, skimmer, ring decelerator, time-of-flight mass spectrometer (TOF-MS), detection laser, and microchannel plate detector (MCP). Not all decelerator rings are depicted. The full decelerator contains 624 rings. The longitudinal stainless steel mounting rods are also not shown. The entire system is contained inside a differentially pumped vacuum chamber.

The voltage sequences for the two deceleration modes create longitudinal potential wells with different ring periodicities and gradients. In the work presented here, both modes of deceleration use maximum voltages of ± 7 kV on the electrodes. PRSD uses four commercial high-voltage switches and TWSD uses eight custom home-made high-voltage amplifiers. Figure 4.2 shows a longitudinal cross section of the ring decelerator in the upper panel and the longitudinal Stark potential for ND_3 for each deceleration mode in the lower panel. The rings are labeled with their rod number, n , from zero through seven. The instantaneous voltages that create the given longitudinal Stark potentials are listed in the table at the top. The left most and right most rings, $n = 4$, are electrically connected, and the voltage and Stark potential pattern is repeated down the length of the decelerator.

The pulsed mode naturally creates wells with greater gradients, which allows it to be more efficient in certain circumstances. While a pulsed-ring configuration with the same well periodicity and gradients as TWSD could be configured, it would be less efficient than the pulsed version presented here. Such a configuration would use the same instantaneous voltages as listed for V_{TWSD} in Fig. 4.2 and create the same instantaneous potentials. The bottom of the Stark potential well would shift by two rings each time and there would only be half the number of switching stages. This reduction in deceleration stages, along with the reduced gradient, greatly diminishes the decelerating capabilities of such a configuration and thus it is not used here.

4.1.2.1 Deceleration Modes

For the PRSD configuration, each rod is connected to one of four Behlke HTS 201-03-GSM-HFB high-voltage switches. These switches are used to produce two voltage configurations that alternate in time. The two voltage configurations produce two Stark potential well configurations, one of which is depicted in the bottom panel of Fig. 4.2. A detailed description of the electrode and voltage configuration can be found in [141]. The time interval between switching from one configuration to another is dictated by the velocity of the molecular beam. These voltages produce an on-axis longitudinal Stark potential with the bottom of the well shifted by l between the two

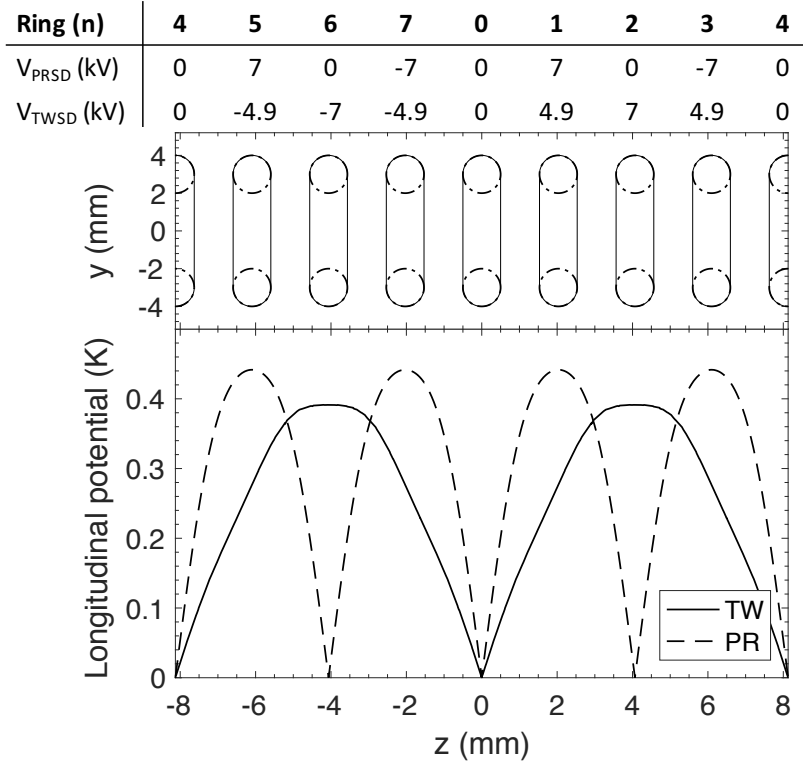


Figure 4.2: Ring number, applied voltages, positions of electrodes, and longitudinal Stark potentials of ND_3 in units of Kelvin. The table and top panel show the ring number n and instantaneous voltages for both deceleration modes plotted below and the longitudinal cross section of the decelerator rings. The bottom panel shows the TWSD and PRSD Stark potentials along the center axis of the ring decelerator. Both Stark potentials are shown at a time when the bottom of the potential well is centered at a ring.

configurations, similar to the pin decelerator. The timing sequence for switching the voltages on the electrodes is calculated in a similar fashion, using phase angles, as that of a pulsed-pin Stark decelerator [114].

For the TWSD configuration, eight analog high-voltage amplifiers are used, one for each rod. These supply the rings with chirped sine wave voltages, which have a fixed phase offset relative to one another. The amplifiers (gain = 12,000) are driven by an arbitrary waveform generator (GaGe CompuGen 8152). The high-voltage linear amplifiers used in this work were designed and constructed in-house. They were tested up to ± 10 kV and in the frequency range of 1-30 kHz on the bench, but were operated at ± 7 kV and in the frequency range of 13-25.5 kHz for this work. These specifications correspond to a molecular beam with velocities from 415 m/s down to 210

m/s. While the amplifiers have a bandwidth of 30 kHz down to DC, the full bandwidth was not utilized since the lower maximum voltage on the rings results in a shallower molecule trap. The set length of our decelerator and this shallower trap did not allow for deceleration down to rest for this particular experiment. However, with ± 10 kV outputs from the amplifiers, the ND_3 molecules with an initial velocity of 415 m/s could be brought to rest.

Each amplifier applies the chirped sine wave voltage V_n , for a given rod n , of

$$V_n(t) = V_a \sin \left(\frac{2\pi n}{8} - \frac{2\pi}{8l} \left(v_i t - \frac{1}{2} a t^2 \right) \right), \quad (4.1)$$

where V_a is the maximum output voltage, v_i is the initial velocity of the molecular beam, a is the acceleration applied by the decelerator, and the $\frac{2\pi n}{8}$ term is the phase shift between rods. The acceleration is given by

$$a = \frac{v_f^2 - v_i^2}{2L}, \quad (4.2)$$

where v_f is the final velocity of the molecular beam and L is the overall length of the decelerator. It is important to note that one period of the sine wave extends over eight ring electrodes and corresponds to two Stark potential wells (see Fig. 4.2). This results from the Stark potential energy depending on the magnitude of the electric field and not the direction.

4.1.3 Simple model of deceleration

The Stark potential wells for each deceleration method give a basic picture into which molecules within the initial phase space can be decelerated. Figure 4.3 shows the transverse potentials for both TWSD and PRSD for the voltage configuration in Fig. 4.2 taken at $z=0$ mm, the bottom of the longitudinal well for both modes. The transverse potential is identical in the x and y directions due to the cylindrical symmetry of the decelerator. The lower panel of Fig. 4.2 and Fig. 4.3 show that for the same peak voltage PRSD has a larger potential well gradient in both the longitudinal and transverse directions due to the reduced ring periodicity of the configuration.

The larger potential well gradient is particularly important in the longitudinal direction. To predict which molecules will be successfully decelerated (i.e. within the phase-space acceptance),




Figure 4.3: The instantaneous TWSD and PRSD transverse Stark potentials in the x-direction for the voltage configuration in Fig. 4.2 and at $z=0$ mm.

we consider the effects of the acceleration on the Stark potential well. Deceleration of the Stark potential well is the equivalent of adding a fictitious potential $U_f = maz$ to the longitudinal Stark potential in Fig. 4.2, where m is the mass of the molecule and z is the longitudinal distance from the center of the Stark potential well. For both modes, the average acceleration is determined only by decelerator length, v_i , and v_f . The acceleration is applied constantly in TWSD, but is time-averaged in PRSD, where instantaneous accelerations are higher. Fig. 4.4 shows the central TWSD and PRSD Stark potential wells from Fig. 4.2 with an acceleration of 0 km/s^2 , -32.5 km/s^2 , and -47.1 km/s^2 added to the Stark potential. For a beam with $v_i = 415 \text{ m/s}$, this corresponds to $v_f = 415 \text{ m/s}$, $v_f = 300 \text{ m/s}$ and $v_f = 230 \text{ m/s}$ respectively. This additional fictitious potential lowers the downstream potential well wall since a is negative, and the effect is referred to as the potential well tilt. (The tilt model of a Stark decelerator is discussed more in depth in [173].) As a increases, the well tips over more and cannot contain and decelerate as many molecules. An acceleration of -47.1 km/s^2 shows how the TWSD Stark potential well is aggressively tilted to the point of almost having no trap depth. Too much acceleration, and thus too much tilt, leads to a small phase-space acceptance and deceleration efficiency.

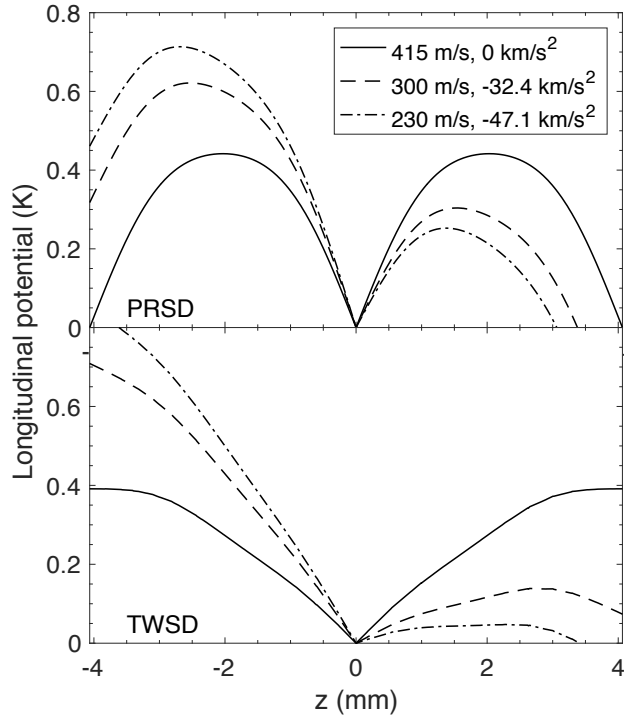


Figure 4.4: The PRSD (top) and TWSD (bottom) Stark potential wells with applied linear fictitious potentials for $a = 0 \text{ km/s}^2$ (solid), $a = -32.5 \text{ km/s}^2$ (dashed), and $a = -47.1 \text{ km/s}^2$ (dot dashed). For a beam with $v_i = 415 \text{ m/s}$, this corresponds to $v_f = 300 \text{ m/s}$ and $v_f = 230 \text{ m/s}$ respectively.

For either mode, too large of an acceleration means no molecules are phase stable inside the Stark potential well. The maximum acceleration where the phase-space acceptance is finite is set by the gradient and height of the Stark potential well. A larger Stark potential gradient allows for more acceleration before the well can no longer hold molecules. In Fig. 4.4 for $a = -47.1 \text{ km/s}^2$, PRSD still has significant well depth, while TWSD has almost none. This is because PRSD has a larger Stark potential gradient. The Stark potential well gradient is set by the deceleration mode and physical decelerator geometry. Decelerating to low final velocities requires either a low initial velocity, which is set by the gas used in the supersonic expansion, or a long decelerator. For the same V_a on the electrodes, the pulsed configuration produces a Stark potential well with a larger gradient in both the longitudinal and transverse directions and a smaller volume because the well only spans $2l$ while the traveling-wave well spans $4l$. This allows the PRSD well to tilt more and be stable at larger accelerations than the TWSD well even though it has a smaller volume.

4.1.4 Simulations

In addition to experimental results, we use Monte Carlo simulations to help interpret and illustrate the differences between the two operating modes. These semi-classical molecular trajectory simulations allow us to explore a wide parameter space that is not yet accessible in the current experimental setup. The simulation starts by creating a packet of molecules at the location of the pulsed valve. The molecules have a Gaussian distribution of initial positions and velocities. The molecular packet is allowed to propagate through a region with no electric field to the entrance of the decelerator. Once the molecules enter the decelerator, their trajectories are determined using a 3D matrix of accelerations derived from Stark potentials modeled by the commercial finite element solver COMSOL. Once the packet of molecules reaches the end of the decelerator both the phase-space distribution (observation at a set time) and the time-of-flight information (observation at a set z -position) are recorded.

Typical initial packets contain between 1 and 8 million molecules that are sampled from a distribution that has full-width half maximum of $\Delta x = \Delta y = 3$ mm, $\Delta z = 30$ mm, $\Delta V_x = \Delta V_y = 5$ m/s, and $\Delta V_z = 30$ m/s. This packet extends beyond the phase-space acceptance of the decelerator.

Since the Stark potential well shapes for both deceleration modes are spatially periodic, the electric field models for determining the molecule accelerations inside the decelerator do not have to span the entire length of the decelerator. Instead, a modular approach can be taken. The position dependent molecular accelerations are calculated from electric field models for a unit cell of 3 (9) rings for PRSD (TWSD). Then, a shift of the longitudinal coordinates within the unit cell allows for molecules to propagate along the decelerator without a model of the entire decelerator. The number of position dependent 3D acceleration models required depends on the operating mode. For PRSD, there are only two voltage configurations and the shape of the Stark potential well is the same for both configurations, but shifted by l . Therefore, only one acceleration model is required and is longitudinally shifted by l each time the high voltage is switched. For TWSD, there is a continuum of voltage configurations corresponding to the sinusoidally varying voltages on the rods.

This is modeled using 25 time steps between the peak of the sine wave being on rod n to the peak being on the subsequent rod $n + 1$. Once the simulations have cycled through all 25 modeled time steps the entire model is longitudinally shifted by l .

4.1.5 Experimental Results and Analysis

4.1.5.1 Deceleration at 7kV

We performed experimental measurements for both PRSD and TWSD at ± 7 kV using the setup described in Sec. 4.1.2 and compared them to simulations. All experiments use a molecular beam with an initial velocity of 415 m/s. The beam was decelerated to various final velocities. Figure 4.5 (4.6) shows PRSD (TWSD) signals for decelerated ND₃ packets as a function of arrival time for selected final velocities. Both figures have results from the simulated decelerated packets on the top and the experimental data on the bottom. There are more decelerated peaks in each packet than shown, but, for clarity, only the central and a couple adjoining peaks are displayed. Both figures have the same vertical scale. The arrival time between peaks in PRSD is more closely spaced than in TWSD due to the well spanning only $2l$ instead of $4l$. Since there is not enough deceleration to separate the decelerated packets completely from the initial pulse, the baselines of the signals are not zero and sit on top of the low velocity tail of the initial molecular beam.

With the same initial velocity, $v_i=415$ m/s, and large accelerations such that $v_f < 220$ m/s, TWSD no longer results in observable decelerated signal, but PRSD still produces clear peaks. PRSD produces observable decelerated molecules down to $v_f = 150$ m/s. While TWSD produces more decelerated ND₃ than PRSD at $v_f > 230$ m/s, it does not produce clearly decelerated packets over the same large velocity range as PRSD, and the amplitude of decelerated peaks decreases more rapidly than in PRSD as the acceleration increases.

The simulations accurately predict the shape and arrival times of the experimental results. While the amplitude of the experimental results and simulation are comparable at higher final velocities (small acceleration), the experimental signal is lower than the simulation at lower final




Figure 4.5: Time-of-flight traces of a molecular beam decelerated with PRSD from an initial velocity of 415 m/s down to the labeled final velocities using the indicated accelerations. The three slowest packet signals have been multiplied by a factor of 4. The top row shows the results of the simulations and the bottom row shows the experimental measurements. The results of the simulations have been vertically offset for clarity.

velocity (large acceleration). These deviations at large accelerations are likely due to imperfections in electrode rings and voltages. This may arise from individual ring misalignment or, for TWSD, deviations in amplifier output voltage from an ideal sine wave. While careful frequency tuning and calibration can mitigate amplifier fluctuations, any imperfections add jitter to the Stark potential well. At large accelerations, jitter may distort the already shallow Stark potential well to the point that many normally phase-stable molecules fall out. However, this would be less of a concern for a deeper potential well with less tilt such that imperfections in the well are small compared to the downstream Stark potential well height of the tilted well.

We use the integrated signal above the baseline in the central peak of each packet to determine and quantify the trade-off between the two modes. The integrated signal for both experimental and simulated results is shown in Fig. 4.7. The x-axis has been changed from final velocity to acceleration using Eq. (4.2), so that the results can be applied to different initial velocities. Now, the




Figure 4.6: Time-of-flight traces of a molecular beam decelerated with TWSD from an initial velocity of 415 m/s down to the labeled final velocities using the indicated accelerations. The two slowest packet signals have been multiplied by a factor of 2. The top row shows the results of simulations and the bottom row shows the experimental measurements. The results of the simulations have been vertically offset for clarity.

trade-off between the two modes is clear. There is a cross-over point between the decelerated signals in TWSD and PRSD. For low accelerations TWSD yields higher overall signal, but at accelerations greater than -46 km/s^2 PRSD yields higher signal. Additionally, PRSD has detectable decelerated packets at accelerations above -50 km/s^2 while TWSD does not. Therefore, if the desired result is a very low velocity molecular beam, PRSD is preferable for this given electrode voltage and configuration.

4.1.5.2 Phase-Space Acceptance

Another method of examining the cross-over point between the two deceleration modes is to examine the phase-space acceptance, the range of initial molecule positions and velocities that will




Figure 4.7: ND_3 signal versus acceleration for both TWSD (yellow squares) and PRSD (blue triangles). The experimental measurements are given by individual points and the results from Monte Carlo simulations by a line. The cross over from TWSD to PRSD producing more decelerated molecule signal occurs near 46 km/s^2 .

be decelerated. We define a cross-over point as the acceleration value where at higher accelerations PRSD begins to produce more decelerated signal than TWSD. One-dimensional (1D) longitudinal (z -axis) phase-space boundary calculations, done in the same manor as [114], are useful for determining the separatrix, which is the boundary between stable and unstable molecular orbits in phase space. Figure 4.8 shows two-dimensional phase-space histograms of all molecules in the central decelerated peak and longitudinal phase-space separatrices for PRSD and TWSD above and below the cross-over point. The phase-space distribution has $z = 0$ as the entrance of the decelerator. Only molecules from the initial distribution that are successfully decelerated and in the central peak are plotted; all other molecules are not shown.

Both modes of deceleration have well-filled phase spaces in both the longitudinal and transverse directions, which differs from the structured phase spaces that are seen in PPSD [117, 139]. Thus, the trade off in total decelerated signal comes from the change in phase-space area and not phase-space filling. From a final velocity of 320 m/s to 200 m/s the phase-space acceptance area of TWSD decreases significantly more than in PRSD. This is due to the significant decrease in Stark

potential well depth from tipping due to the acceleration. While the longitudinal phase-space acceptance captures the phase-space boundary of the decelerated molecules, the area of the phase space does not accurately model the efficiency of the deceleration at different final velocities. Even a 3D phase-space volume calculation fails to model the rate of decrease in experimental signal with increased acceleration. This is because the total number of successfully decelerated molecules is determined by the overlap of the phase-space acceptance with the initial distribution of molecules. To account for the initial distribution of the beam, the deceleration efficiency should be computed using Monte Carlo simulations.




Figure 4.8: Longitudinal phase-space plots with TWSD in the top row and PRSD in the bottom row. The left column has an acceleration of -27.6 km/s^2 and the right column has an acceleration of -52.2 km/s^2 . For an initial velocity of 415 m/s, this would be a final velocity of 320 m/s and 200 m/s respectively. The solid-red lines are the 1-dimensional separatrices. There is no significant structure in the phase-space acceptance in either deceleration mode.

4.1.5.3 Determining the Optimal Deceleration Mode

When using decelerated molecules in an experiment, the deceleration mode would ideally be chosen to have the maximum decelerated signal for a given final velocity. This may not always be an option due to hardware constraints. However, even if the choice is available, it is not always clear which mode is more efficient. The trade-off between the two deceleration modes is especially

relevant given experimental limitations of a given decelerator. To determine which mode is best to use, different voltage regimes were explored using simulations. The same decelerator geometry as previously described was used while changing the peak voltage in both modes. Results of these studies are shown in Fig. 4.9 for 5, 7, and 9 kV. A crossing point where TWSD and PRSD produce the same signal exists for these voltages and occurs at a larger acceleration for greater electrode voltages.




Figure 4.9: Simulated integrated ND_3 signal versus acceleration for both TWSD (squares) and PRSD (triangles) with peak voltages of 5, 7, and 9 kV. A crossing point exists for all voltages shown.

Generally, TWSD is more efficient for mild accelerations, but for large accelerations, especially at low peak voltages, PRSD may be more efficient. The crossing point may or may not be relevant depending on the parameters of the initial beam and maximum voltage on the electrodes. For example, while a crossing point for 10 kV (not plotted in Fig. 4.9) exists, it is not accessible in our decelerator using a molecular beam seeded in krypton. The acceleration to bring ND_3 seeded in krypton from 415 m/s down to rest is -67.9 km/s^2 and the crossing point at 10 kV is near -74.3 km/s^2 . Thus, for 10 kV on the electrodes, a beam seeded in krypton, and this decelerator geometry, TWSD will always produce more ND_3 at all final velocities.

However, at lower peak voltages, PRSD is an important operating mode to consider. The advantage of the PRSD is that it allows access to higher accelerations than TWSD due to the greater Stark potential well gradient. As controlling high-voltages is one the largest experimental challenges for any Stark deceleration experiment, PRSD may be an important tool to investigate low final molecular velocities given tight voltage constraints. This is because PRSD produces more signal at higher accelerations and lower electrode voltages and is easier to implement than TWSD. PRSD also allows the beam to reach a lower final velocity, possibly one that can be trappable, without having to build a very long decelerator or obtain analog high-voltage amplifiers.

For a given decelerator length, a quick method of determining which mode will yield more decelerated molecule signal is desirable. Previously, the crossing points between the two modes were determined by integrating time-of-flight traces from Monte Carlo simulations, which is computationally intensive. Instead, the crossing point can be determined using an efficient 1D separatrix calculation.

The 1D separatrix area does not accurately model the decrease in experimental signal with increasing acceleration for either model since it does not account for the initial molecule position and velocity spread. For the separatrix area to accurately predict the rate of decrease of the decelerated peaks the overlap of the initial distribution and the phase space must be known. However, this overlap function affects the phase space of both modes similarly and does not change the point where their phase-space areas are equal. Thus, calculating the acceleration when the longitudinal phase-space areas for both modes is equal gives a good estimate for where the crossing-point will occur.

Figure 4.10 shows the crossing points determined by both Monte Carlo simulations and 1D separatrix calculations are linear with respect to applied acceleration. The experimentally determined crossing point, using the data from Fig. 4.7, is plotted with an open circle and agrees with both calculations. The 1D separatrix calculations agree with the Monte Carlo simulations and is a simple method for determining the crossing point, even if it does not quantify how much better one mode is compared to the other. This simple 1D separatrix calculation easily predicts when one

deceleration mode is more efficient than another without having to run Monte Carlo simulations.




Figure 4.10: The acceleration at the crossing-point, where TWSD and PRSD produce similar decelerated molecule signal, calculated using Monte Carlo simulations and 1D longitudinal phase-space acceptance. The crossing point determined by either method is linear with respect to applied voltage. The experimental crossing-point (open circle) agrees well with the calculations and simulations.

4.1.6 Conclusion

Stark deceleration has proved to be a reliable source of cold molecules for various applications. However, traditional crossed-pin geometry Stark decelerators have well documented instabilities that reduce their efficiency. An alternative deceleration technology, traveling-wave Stark deceleration, is predicted to nearly eliminate the problems that reduce the number of decelerated molecules. However, the difficulty in creating the required sinusoidally varying high-voltage waveforms for TWSD has limited their integration into experiments. To possibly make implementation of a ring-geometry decelerator less challenging, we investigated running this type of decelerator in a pulsed mode.

We demonstrated using experimental measurements and molecular trajectory simulations that for the same peak voltage on the ring electrodes, the alternative running mode, PRSD, is more

efficient than TWSD at high accelerations. This effect is caused by the larger Stark potential gradient for PRSD, which allows for larger accelerations before the stable phase-space region vanishes. The crossing point at which the two modes are equally efficient depends on the maximum applied voltage, among other experimental parameters. The PRSD mode is a particularly important option for use with lower applied voltages, and gives access to a previously inaccessible low velocity regime. Thus PRSD provides a useful and straight forward method for decelerating many molecular species in a ring-geometry decelerator. These results could allow for more groups to be able to implement ring-geometry Stark deceleration using the less challenging electronics, and still gain many benefits from the cylindrical symmetry of the electrodes.

4.2 Implementing PRSD




Figure 4.11: A schematic diagram of a cross section of the ring decelerator. The two voltage configurations above and below give the two different potentials used in PRSD. The numbers on each ring correspond to the rod that it is connected to. Thus the left-most and right-most rings are both connected to rod 7.

The high-voltage electronic configuration in PRSD is similar to that of PPSD, but instead of four rods that each must be switch between +HV or -HV and ground, there are eight. Two alternating voltage configurations are used to create the potentials wells, Figure 4.11. Each ring is labeled with its rod number, and the two possible voltage configurations are given above and below the rings. Although there are more rods in the ring decelerator geometry, only four high-voltage

switches are required since every fourth ring has the same voltage sequence. This results in two Stark potential wells formed for every eight rings (Figure 4.2). Here, it is again important to note that the Stark potential well is only half the longitudinal length in PRSD compared to TWSD. This will play a significant role in the cases where PRSD is more efficient than TWSD.

The similarities between PPSD and PRSD mean the digital timing sequences for switching between the two sets of potentials are calculated in the same manner (see Section 2.2). One timing molecule located at the center of the initial pulse and with no transverse displacement or velocity is propagated along the on-axis longitudinal potential at a given phase angle. In addition to the same method of calculating the timing sequence, the physical equipment for implementing PRSD is straightforward as it uses essentially the same electronics as PPSD. A few modifications to the electronics maybe necessary due to the longer digital timing sequences in PRSD and the increased capacitance of the decelerator. The number of digital voltage switches times required will be at least the number of deceleration stages. For our pin decelerator, this is 149 if operated in $S=1$ mode. For our ring decelerator this is 623. This is one less than the number of pin pairs/rings in the decelerator since each deceleration sequence begins with the timing molecule in between the first and second pin pair/ring. The number of deceleration stages is then equal to $n - 1$ where n is the number of pin pair or rings.

Just like in the case of PPSD, the high voltage potentials are switched using Behlke HTS-201-03-GM switches. In our PPSD, we used a Viewpoint DIO64 board to generate the control sequences for the Behlke switches, but the DIO64 cannot be used for the digital timing sequence in PRSD. It does not have sufficient memory to generated the required timing sequence for PRSD, which for our decelerator requires over 623 logical switches. At least one voltage switch time is needed per stage and a few more are needed for the valve and detection sequence. The DIO64 has a maximum of 512 lines of output making it insufficient for this application. Instead, we use a GaGe 8152 arbitrary waveform generator, which we also use in TWSD, to create the switching sequence. It has a memory length of 512K samples per channel and is typically run at 37.5 MHz for a maximum output sequence length of 13.65 ms. Since the GaGe board only has a maximum

output voltage of ± 1 V, a 1V to 5V amplifier is required in order to convert the output signal to a TTL to drive the Behlke switch control circuitry. We accomplish this with a Schmit trigger using an AD790 comparator.

We use four Behlke switches to control the eight rods. An additional high-voltage cable splitter box is added to each switch output since each high-voltage switch must drive two decelerator rods instead of one. The two voltage configurations each require a different timing control sequences. Additionally, for each timing control sequence, one active high version is required for the positive voltages and one active low version is required for the negative voltages. The GaGe board creates the active high version, for +HV, of both timing control sequences and then an inverter box produces the active low sequences, for -HV, from the active high sequence. The inverter box also ensures that the control signal for the active low sequence always remains high when the decelerator rings are not suppose to be switching so that the output remains at ground. While all four switching sequences could be programed into the GaGe board, this configuration reduces the required number of channels required.

The +5 V digital timing sequence is passed through an optoisolator before going to the Behlke control circuit. This protects the computer from high-voltage switching noise or failures. The four digital timing sequences drive four Behlke switches, which operate in the high frequency burst mode. The maximum burst frequency for our switch model is $f_{v(max)} = 3.3$ MHz. The shortest switch time step, which occurs in bunching mode, is $4.8 \mu s$, which is well within the specifications of the Behlke switch. However, since there are many more switches per burst required in PRSD than in PPSD, it is important that the Behlke switch control box have enough capacitance to drive all the pulses. The recommended Behlke control circuit is shown in Figure 4.12. For the burst option, the capacitance of C2 and C3 is recommended to be

$$C = 0.1 \mu F \times \text{number of pulses per burst.} \quad (4.3)$$

Due to space constraints, we use $56 \mu F$ capacitors, as larger capacitors do not fit into the control circuit board. We found that $56 \mu F$ is sufficient for our switching sequence even though it is less

than suggested by Equation 4.3. Here, it is also important to note that in our circuit we use a 25 μF capacitor for the +5 V supply backing capacitor in place of the 10 μF capacitor between pins 3 and 5.

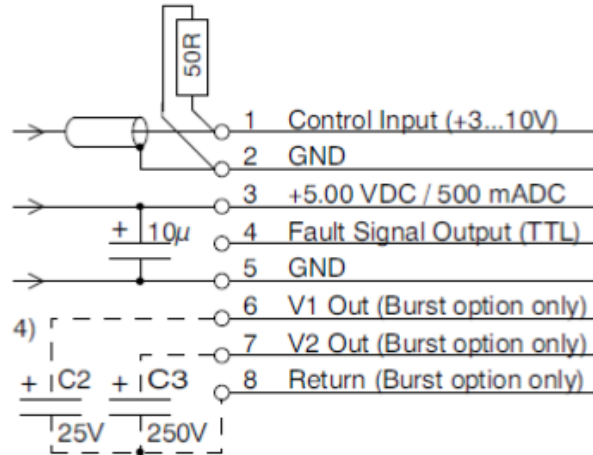


Figure 4.12: Recommended Behlke control circuit from Behlke HTS-201-03-GSM technical data sheet. For the burst option the capacitance is recommended to be $C=0.1\mu\text{F} \times \text{number of pulses per burst}$.

The last modification required to switch from PPSD to PRSD is on the output of the Behlke switch. Since each switch is connected to two rods, this increases the capacitance that each switch must drive. In order to speed up the switch output, we removed the typical 1.1 k Ω resistor on the output of the Behlke switch. This normally protects the switch in an over current event such as an arc. Rise times for various Behlke switch output resistor and decelerator rod or test capacitor configurations are given in Table 4.1. The rise time for a square wave on the pin decelerator is 0.8 μs . In order to achieve a similar rise times to that of the pin decelerator, while still having each switch connect to two rods, the output resistor must be removed.

Table 4.1: Switching rise times for PPSD vs PRSD

Load	Output resistance (Ω)	Rise time (μs)
Pin decelerator, 1 rod	550 or 1.1 K	0.8
Ring, 2 rod	1.1 K	1.6
Ring, 2 rod	550	1.3
Ring, 1 rod	1.1 K	0.8
250 pF test load	0	0.7

4.3 Decelerating with PRSD

While the comparisons between PRSD and TWSD in section 4.1 were only performed up to ± 7 kV, we continued PRSD to higher voltages. At ± 7.5 kV on the decelerator, very clear decelerated peaks with adequate signal-to-noise can be seen for final velocities down to 75 m/s, Figure 4.13. We did not pursue lower velocities or higher voltages than this due to an issue with the decelerator. With ± 7.5 kV on the electrodes, the decelerated signal would occasionally drop by a factor of 2-3. This drop could not be attributed to laser or valve behavior, and we believe it was an issue with the potentials on decelerator rings even though the decelerator had been conditioned in pulsed mode to higher voltages. Switching the polarities on the decelerator rings recovered the signal. However, deceleration at ± 7.5 kV was not continued since it appears that further voltage conditioning of the decelerator is required.

As mentioned in Section 4.1.5.2, the 1D phase-space model does not accurately depict the functional form of the decrease in molecule signal as a function of acceleration. The 1D LPSA models can only be interpreted qualitatively and not quantitatively. The relative ratios between the LPSA at different accelerations is not directly proportional to observed decelerated signal since each phase-space area must also be convolved with the molecule density. However, 1D longitudinal phase-space acceptance plots are still useful for visualizing the parameter space we are working with. Here, we take the 1D model for the lowest voltage and velocity configuration that we can detect in the lab as a limit of what it is possible to decelerate. Any voltage and velocity combination that has a 1D phase-space acceptance larger than that limit can be decelerated and detected in




Figure 4.13: PRSD of ND_3 from an initial velocity of 415 m/s down to final velocities of 125, 100, and 75 m/s.




Figure 4.14: Longitudinal phase-space acceptance for PRSD of ND_3 from an initial velocity of 415 m/s down to various final velocities using peak voltages of ± 7 , 7.5, and 8 kV.

our decelerator. This simple model makes the assumption that molecular dynamics are not a function of phase-space acceptance. 3D molecular dynamics simulations should still be conducted before drawing any conclusions. However, this model serves as a quick method to estimate what is

achievable at different voltages.

Figure 4.14 shows the 1D phase-space acceptance for various final velocities and peak voltages. The LPSA for the two lowest velocity configurations that we have detected in the lab, ± 7 kV, $v_f=150$ m/s (Figure 4.5), and ± 7.5 kV, $v_f=75$ m/s (Figure 4.13), are shown in red and orange respectively. The smallest phase-space acceptance of these two is for ± 7 kV and $v_f=150$ m/s. Any LPSA that has an area larger than this case should have signal with comparable or better signal-to-noise than the blue trace in Figure 4.5. For ± 7.5 kV deceleration with half the final velocity (orange), the LPSA is larger than the ± 7 kV case, showing the importance of the well depth and gradient. The green trace shows the LPSA for ± 7.5 kV deceleration down to $v_f=25$ m/s, a trappable velocity. Since the acceleration required to reach a final velocity of 25 m/s does not differ much from that required to reach 0 m/s, the LPSA does not change significantly for $v_f < 25$ m/s. This LPSA is only slightly smaller than the ± 7.5 kV, $v_f=75$ m/s case and larger than the ± 7 kV, $v_f=150$ m/s case. This suggests that in our particular decelerator, we can tune over the full velocity range using ± 7.5 kV PRSD. Higher voltages from here will increase the well depth, which results in greater phase-space acceptance and more decelerated molecules. When the peak voltage is increased to ± 8 kV, all LPSAs are larger than the slowest cases that we have decelerated experimentally. The ± 8 kV LPSA in Figure 4.14 (blue) is the largest acceptance and once again the LPSA for $v_f < 25$ m/s is approximately the same as that for $v_f = 25$ m/s. Thus, if one did not want to use high-voltage amplifiers and TWSD, this decelerator (624 rings, 2.032 mm apart) in PRSD mode would be able to decelerate ND_3 to any desired velocity starting at voltages of ± 7.5 kV.

Additionally 1D phase-space acceptance model can be useful for visualizing when to use each deceleration technique with various carrier gases. Figure 4.15 shows the 1D model for ± 7 kV deceleration. Vertical lines have been placed at the accelerations required to bring an expansion of ND_3 in krypton and xenon to rest. The left side of each vertical line represents the accessible range for that gas. Here, we see that if working with an expansion in xenon, TWSD will always be superior at this voltage. PRSD only shows superior performance when more deceleration is

required (*ie*, in a faster krypton beam).

The same style plot is shown in Figure 4.16 for ± 10 kV PRSD. For all accelerations that would be used when working with a beam seeded in krypton, TWSD is superior. Unless one wanted to work with argon or did not have access to TWSD, PRSD at ± 10 kV is not as useful.




Figure 4.15: The LPSA for ± 7 kV PRSD as a function of acceleration. Vertical dashed lines indicate the acceleration required to decelerate ND_3 seeded in xenon and krypton down to rest.




Figure 4.16: The LPSA for ± 10 kV PRSD as a function of acceleration. Vertical dashed lines indicate the acceleration required to decelerate ND_3 seeded in xenon, krypton, and argon down to rest.

4.4 Transverse Dynamics

The transverse potential energy landscape between the two deceleration modes evolves differently over time. However, this does not appear to result in significant differences in the final decelerated packet. Just like in the longitudinal picture, the well depth is the primary determinant of the phase space acceptance. Figure 4.17 shows the Stark Potential energy surface in the xz ($y=0$) plane for both PRSD and TWSD. Both plots are taken at times where the bottom of the potential well is aligned with a ring electrode located at $z=0$ mm. Since the decelerator is cylindrically symmetric, the potential energy surface looks the same in the xz ($y=0$) and yz ($x=0$) planes.

Although the two surfaces look similar, except for twice as many wells longitudinally in PRSD compared to TWSD, in this snapshot in time, they evolve rather differently in time. In TWSD, potential well bottom is always moving and therefore has an infinite number of configurations. Luckily, when the bottom of the potential well is not aligned with a ring, the shape of the well only changes slightly (Figure 4.18). Here z is given in units of L , the ring spacing, since fractions of 2.032 mm are difficult to read and remember. The major deviations of the transverse potential




Figure 4.17: The Stark Potential energy surface in the $xz(y=0)$ plane for both PRSD (left) and TWSD (right). Both plots are snapshots for when the bottom of the potential well is aligned with a ring located at $z=0$. *need to add a , b labels to plots*

occur at $|x| > 1$ mm, which is significantly displaced from the bottom of the trap. Figure 4.18b shows a closer view of the bottom of the transverse potential at various locations. While there is variation, it is small compared to the overall trap depth.

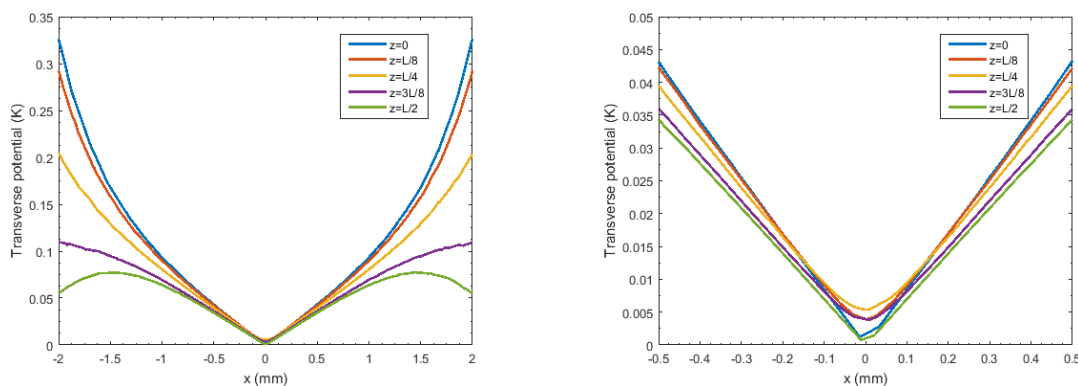


Figure 4.18: The transverse Stark potential for ND_3 in the TWSD when the bottom of the well is ring aligned ($z=0$) and gap aligned ($z=L/2$). The plot on the right is a closer examination of how the bottom of the potential well varies as it moves from being ring aligned to gap aligned. *Maybe make this an inset instead*

In PRSD on the other hand, the electric fields remain static for longer and change discretely as opposed to continuously. In this case, the timing molecule traverses the potential energy surface instead of always residing at the bottom. For example in bunching, the timing molecule starts part way up the side of the potential well, it travels through the bottom of the potential well and up to

the same height that it started at. As a result, it encounters a transverse potential that changes more drastically than in TWSD. The transverse potential encountered by the timing molecule for the pulsed-ring case can be thought of as taking successive slices of the 3D potential well in the radial direction. In the traveling-wave case, on the other hand, the entire potential well translates with the timing molecule and the transverse potential does not significantly change as viewed by the timing molecule. Figure 4.19 shows the transverse potential that the timing molecule encounters for both types of deceleration. The left most plot $z = 0$ is where the timing molecule is aligned with a ring. The right most plot is when the timing molecule is aligned with a gap (exactly between two rings). Here, PRSD is taken to be in bunching mode, so the potential shifts by L every time the timing molecule is gap aligned ($z = \frac{L}{2}$). The pronounced different shape between the TWSD and PRSD configurations is due to the significantly deeper well of PRSD. This can also be seen in Figure 4.18 where the PRSD well has much higher walls (yellow on the edges). Despite the differences between the two transverse potential energy landscapes, a preliminary examination of the transverse phase space for both PRSD and TWSD shows that both are evenly filled and no inherent structure is seen in the distribution of molecules. A detailed analysis of the transverse behavior of molecules was not carried out. A more in depth analysis is probably needed to see if weak transverse-longitudinal coupling is present.

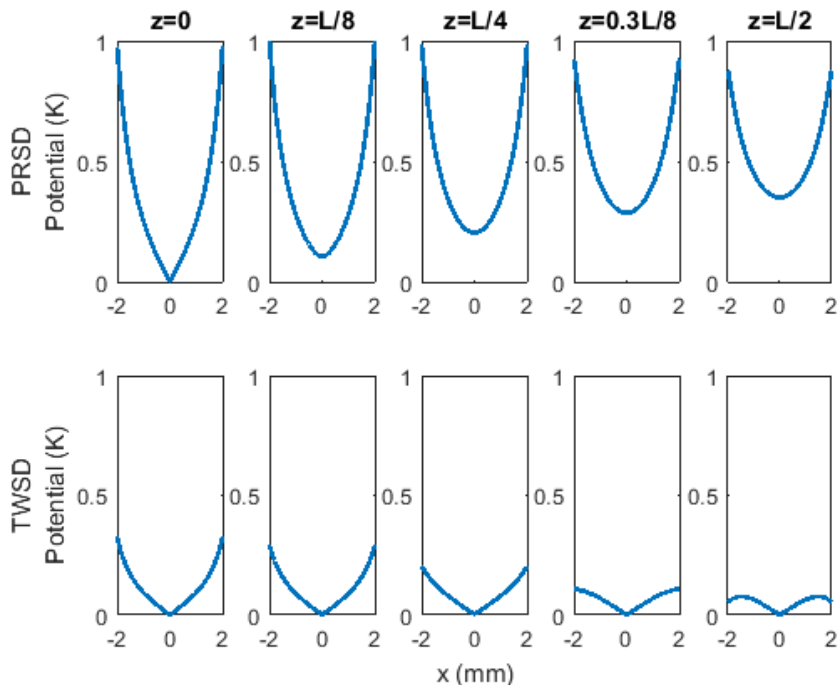


Figure 4.19: The transverse potential that the timing molecule encounters at various point between being ring aligned (left most) and gap aligned (right most). In the PRSD case the potential well does not move while the timing molecule climbs up the potential hill, this results in the bottom of the transverse potential well increase in energy as the molecule gets close to $z = \frac{L}{2}$.

4.5 Calibration

Stark deceleration data such as those in Figure 4.5 and 4.6 must be calibrated before direct comparisons can be made between data from different final velocities and between experiment and simulation. Differences in different sets of experimental data and from experimental data to simulation data must be accounted for. Recording time-of-flight data in deceleration experiments usually required many measurements of the same deceleration sequence in order to have sufficient signal-to-noise. Each data set for a given final velocity (each colored trace in Figure 4.5) is referred to as one scan. The experimental time-of-flight scans presented in Section 4.1 are averages of n runs of the experiment per data point. *figure out how much time is actually*. During this time the experimental signal can drift due to changes in temperature and noise from other equipment. In order to ensure many time-of-flight scan can be compared to each other, a simple calibration

method is used to account for day-to-day signal drifts. While the general envelope of the signal remains the same in day-to-day runs of the experiment, the amplitude can be greatly affected by the precise gas mixture (a 2.25 L, 60 psi gas mixture will last about 100-120 hours of running at 10 Hz) and laser power, wavelength, and beam profile. To account for these drifts during the course of a detailed time-of-flight scan (1-2 hours), we use a simple linear calibration between two sets of points to account for laser wavelengths drifts. All PRSD scans presented in this chapter were taken with a single gas mixture, so in this case we were only calibrating for laser drift.

One scan is typically (*give parameters*). The data points are usually taken sequentially in time, however they could also be randomized in computer control program. When comparing many different time-of-flight traces to one another, such as in Figure 4.5, one time-of-flight scan must be chosen as the ‘master scan’ all other scan will then be calibrated to this one. Taking the scans in Figure 4.5 as an example, if the $v_f=240$ m/s scan (yellow) were chosen as the master scan, all other scans will have their amplitudes calibrated relative to it. The choice of master scan is not critical since decelerated signal is only compared relative to other decelerated signal, however the master scan should be one with adequate signal-to-noise and in no danger of saturating the detection system.

In order to calibrate all scans to the master scan, two calibration points are chosen and measured before and after every scan. The calibration points are usually chosen to be near the peak of one deceleration packet. The two calibration points should be taken from decelerated packets at different final velocities with significantly different signal amplitudes. This way, in the case of large laser instabilities that cause one point to go into saturation or become impractically low, there is a back up calibration point. Before and after each time-of-flight deceleration scan, a measurement of both calibration points is taken. These points are referred to as C_{im} and C_{fm} where m is either 1 or 2. Typically only C_{i1} and C_{f1} are used, but C_{i2} and C_{f2} are recorded as backup calibration points. Each calibration point is made of measurements from at least 200 shots of the experiment. Typically this consist of 64-128 averages on the oscilloscope and 3-8 averages on the computer.

Although the source of the drifts may not be linear, the drift between the two points is assumed to be linear over time. For a time-of-flight scan containing n raw data points y_n , each calibrated point y_{adj_n} is given by

$$y_{adj_n} = y_n \left(1 + \frac{C_{im} - C_{fm}}{C_{im}} \times \frac{t_n - t_i}{t_f - t_i} \right) \frac{C_{pm}}{C_{im}}, \quad (4.4)$$

where

m is either 1 or 2 and refers to whether the first (1) or backup (2) calibration points will be used,

C_{im} is the initial calibration point for the scan which was taken immediately before the scan began,

C_{pm} is the initial calibration point of the “master” scan that all scans are referenced to, the $C_{pm} = C_{im}$ for the master scan

C_{fm} is the final calibration point for the scan which was taken immediately after the scan finished,

t_n is arrival the time for the measured point y_n ,

t_i is the arrival time of the first data point of the given time-of-flight scan, it corresponds to the time at the point y_1 ,

and t_f is the arrival time of the last point in the given time-of-flight scan, it corresponds to the time at points time at $y_{n_{max}}$.

The initial calibration point for the “master” scan is taken as C_p . If calibrating only one scan, then $C_p = C_i$. This is particularly necessarily for very long (several hours) time-of-flight scans where both the underlying undecelerated peak and the decelerated peaks must be recorded if there are laser drifts.

In the case of comparing multiple sets of scans taken under different conditions (deceleration mode, laser power, gas mixture, etc.), the master scans for each set can be calibrated to each other via free flight scans. A free flight (no deceleration) scan should be taken at the beginning and end of each day. The relative peaks of these free flight scans can then be used to scale between two master scans. This allows for comparisons of relative signal, but should only be used if no major adjustments to the system (detection laser position, detection voltages) have been made.

This calibration was found to be reliable when comparing calibrated scans from days when the laser was rather unstable to days where the laser was more stable. By applying this calibration, we saw very reliable signal across many days and even spread out across several months. Figure 4.20 shows the difference between raw uncalibrated data and calibrated data. Here, three different raw time-of-flight scans of ND_3 decelerated from $v_i=415$ m/s down to $v_f=240$ m/s are displayed (Figure 4.20a). These scan were recorded on two different days. The uncalibrated time-of-flight traces have widely varying amplitudes, but once the calibration is applied, all three scans have comparable amplitudes. The calibration takes into account drifts such as the behavior of the pulsed valve and the wavelength of the dye laser. Although these effects may not necessarily have linear drifts, this linear calibration still does a good job of making signal comparable from one day to the next.




Figure 4.20: Three ND_3 PRSD time-of-flight scans for $v_i=415$ m/s down to $v_f=240$ m/s of ND_3 . On the left is the raw scans. On the right after the calibration is applied and the scans trimmed to be the same length in time. The colors remain the same between both plots

Once all experimental scans have been calibrated to each other, they maybe compared to simulations. Since only the relative signal of the experimental scans are being compared, they can be scaled by an arbitrary factor or normalized such that bunched signal (or any other decelerated peak) has a value of one. In the experiment the detected molecule signal is measured in nanovolt-seconds (nVs) since is the integral of a voltage versus time trace. The value of the signal is dependent on the voltage on the MCP detector and gain of the particular detector plates. Thus it should only

be viewed as a relative measure and a constant multiplier can be used to compare experimental scans and simulated scans.

Additionally, due to uncertainty in the free flight distance between the exit of the decelerator and the position of the detection laser, an arrival time adjustment must be applied to the simulation data when comparing it with experimental data. The arrival time shift, t_{shift} , takes the form of

$$t_{shift}(v_f) = \frac{d_0}{v_0} + \frac{d_f}{v_f}, \quad (4.5)$$

where d_0 and d_f are constants that account for differences between the simulation and experiment. The first term $\frac{d_0}{v_0}$ results from uncertainty in the distance between the valve and the entrance of the decelerator. It is a constant time shift regardless of the final velocity of the beam. In the simulations, the timing molecule can automatically be loaded into the correct potential well since all distances can be set exactly. In the experiment, the position between the valve and skimmer can be adjusted to find the best free flight beam. Then the time between the valve firing and the decelerator turning on can be adjusted for optimal beam loading. However, due to the adjustability of the valve position, and the delay between the valve control pulse and when it actually opens, the time it takes a molecule at the densest portion of the molecular packet to travel from the valve face to the decelerator entrance is not precisely known. d_0 accounts for the differences in free flight time between the valve and the entrance of the decelerator. d_f on the other hand, accounts for the distance between the end of the decelerator and the detection laser. Since the arrival time will depend on the final velocity, this must be specified as a distance not a time. While these values will change depending on how close the values used in the simulation are to the experiment, for reference the calibration for PRSD in (Figure 4.5) were $d_0 = 9.46$ cm and $t_0 = -9.27$ mm.

4.6 8-ring periodicity PRSD

The discussion up to this point has been for a voltage configuration where the voltage sequence repeats every four rings, and four potential wells are formed for every eight rings in the decelerator. This is referred to as the 4-ring periodicity pulsed-ring Stark decelerator (4R PRSD). Since there are

8 electrically isolated rods in our ring decelerator, there are other possible pulsed-ring configurations. As alluded to in 4.1.2 an 8-ring periodicity pulsed-ring Stark decelerator (8R PRSD) can also be implemented. This configuration is not as effective since it is more difficult to run and less efficient than the 4R version at high accelerations, but warrants discussions since it creates instantaneous Stark potential wells that are identical to those in TWSD.

An 8R PRSD configuration would have the same voltages as TWSD when the bottom of a well is centered on a ring. Thus ± 7 kV 8R PRSD would have the same potential as TWSD in Figure 4.2 and the same instantaneous voltages. Figure 4.21 shows both of the 8R voltage configurations. In one of the static configurations, every fourth ring is at ground (rings 0 and 4 in the case of the upper configuration). The rings halfway between the grounded rings are at +HV and -HV (rings 2 and 6), and all the remaining rings (all the odd ones) are at 71% of the high voltage. The odd rings are at 71% and not at 50% of the high voltage because the smoothest potential is created by a sine wave. As seen in the TWSD case for a sine wave voltage when the zero crossing is centered at a ring, the intermediate rings are at $\sin\left(\frac{\pi}{4}\right) = 0.707$. If the odd rings were at 50% of the high voltage the top of the Stark potential hill would have a very small local minimum.

In order to run the two configurations shown in 4.21, four power supplies are required and six high-voltage switches. It is interesting to note that in this configuration, rods 1 and 5 never switch voltage and rods 3 and 7 require bipolar switching. While this is experimentally possible, it makes running this type of deceleration much more complex than 4R PRSD.

Figure 4.22 shows the comparatively low efficiency of 8R PRSD using simulated data. The initial velocity for all simulations is $v_i = 415$ m/s and the molecules are decelerated to final velocities of $v_f = 415$, 360, and 320 m/s. The difference in well periodicity is easily seen in the time-of-flight traces with two 4R peaks occurring for every 8R packet. While in bunching ($v_f = 415$) the 8R configuration produces greater time of flight signal, this signal quickly decreases once the Stark potential wells are decelerated. At a final velocity of 360 m/s ($a = -16835$ m/s²), the peak amplitude of the decelerated packets is comparable for both periodicities, and at a final velocity of 320 m/s ($a = -27578$ m/s²) the peak amplitude for 8R is significantly less than that of 4R deceleration.

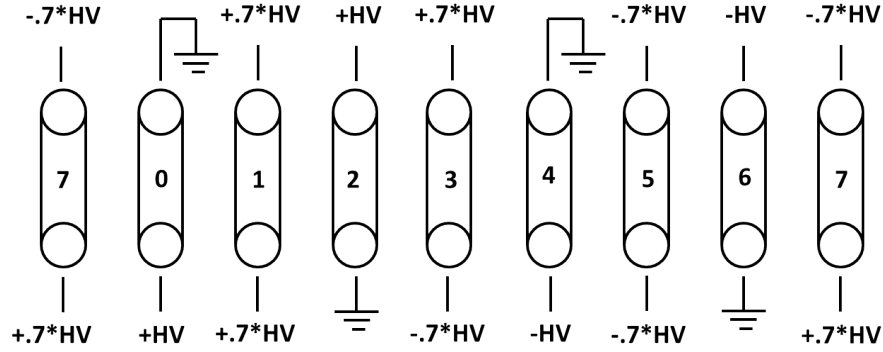


Figure 4.21: The two voltage configurations for 8R PRSD are given above and below the cross section of the ring decelerator. The numbers on each ring correspond to the rod that it is connected to.

Although the Stark potential well in 8R deceleration is larger, which increases the phase-space acceptance in bunching, due to its lower gradient, it tilts over more quickly as acceleration is increased. Additionally, since the bottom of the potential well moves by two rings during each switch, in 8R PRSD, our decelerator has only half as many deceleration stages. It performs worse than 4R PRSD except in cases of very low accelerations and bunching, and it always performs worse than TWSD since the instantaneous potentials are the same, but the well propagates discretely instead of continuously.

While having identically periodic potentials for the pulsed and traveling-wave method would be ideal for comparing the detailed molecular dynamics within each of the two methods, 8R PRSD is not an efficient method of running the decelerator. Instead, 4R PRSD is a straightforward method that can be used by groups who may have limited space or voltage constraints. It is a method that is not more difficult to assemble than PPSD, can produce similar, and in specific circumstances, better results than the much more difficult to implement TWSD method.

Of course there are even more permutations of pulsed configurations, however these become increasingly difficult to implement experimentally. For example, one can imagine a configuration where the bottom of the potential well moves only one ring per switch, just like in 4R PRSD, but now with a potential well that is larger. This would come closer to approximating TWSD but each

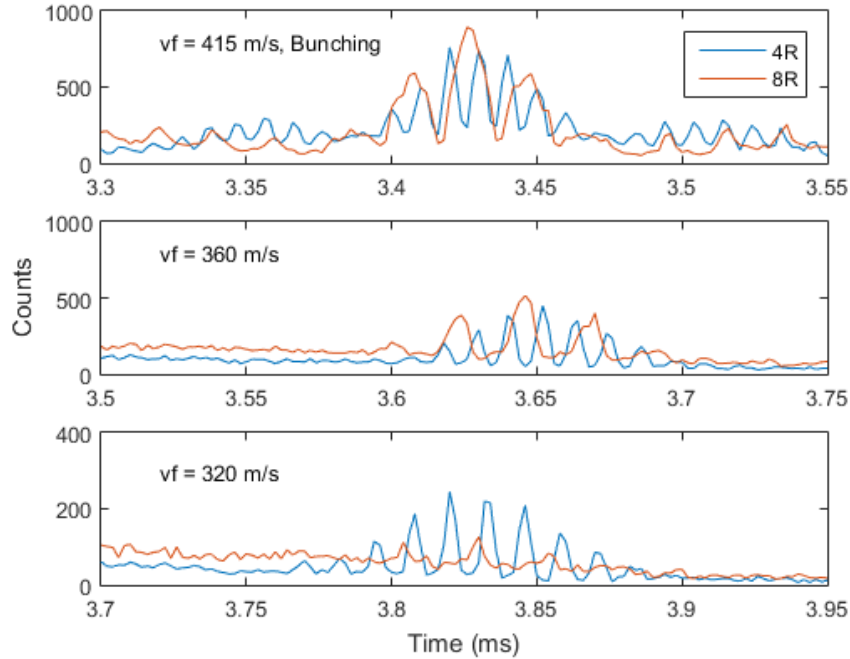


Figure 4.22: Simulated data of 4R and 8R PRSD at ± 7 kV with $v_i = 415$ m/s and $v_f = 415, 360,$ and 320 m/s.

ring would have to switch between three or four voltages and a combination of three and four-state switches would be required. This would also greatly change the phase angle calculations since the well does not move by $\phi = 180^\circ$ each time. The complexity of these cases illustrates why pulsed-ring configurations beside the 4R configuration in Figure 4.11 are unlikely to be used experimentally since, with a subset of the same equipment, it could be operated more efficiently.

Chapter 5

High-Voltage Amplifiers

The most difficult step towards the experimental realization of TWSD is generating the required high-voltage waveforms. As discussed in Chapter 2, very few experiments have used this technique due to the need for stable amplifiers with an output voltage greater than ± 5 kV and a large bandwidth or a long ring-electrode decelerator. Here, we present an amplifier that has been successfully used to decelerate ND_3 with ± 7 kV output voltages and has the capabilities to reach higher output voltages. Individual amplifiers have performed adequately at output voltages up to ± 12 kV, and with a few modifications, these amplifiers will be able to be used for TWSD of a ND_3 beam down to trappable velocities.

The first section of this chapter is adapted from a manuscript under review [174]. It discusses the amplifiers in detail and contains circuit diagrams for important portions of the amplifier circuit. This is followed by a discussion of the control circuitry and interlock circuitry to illustrate how the amplifiers are integrated into the experiment. For a full in-depth discussion of the amplifiers and all associated circuitry, background on the individual subcircuits, and tuning and testing procedures see Noah Fitch's dissertation [95]. It is important to note that these amplifiers operate at very high voltages and must be carefully insulated; Section 5.6 discusses construction and insulation to keep both the user and the amplifier safe. Several changes have been made to the amplifiers and control circuitry since Noah's dissertation including small changes to resistor values and a few larger design changes. These are covered in Section 5.8 and the full updated circuit diagram for the entire amplifier setup can be found in Appendix A.

5.1 High-Voltage Amplifier Electronics

5.1.1 Introduction

Stark deceleration is a molecular manipulation technique that uses time-varying inhomogeneous electric fields to produce packets of slow, cold molecules [107]. In this technique, large electric-field gradients are created inside a vacuum chamber via application of large potential differences between electrodes. These gradients exert Stark-effect forces on polar molecules in particular quantum states. With the correct electrode geometry and temporal control of the applied potentials, molecular beams can be decelerated. The resulting beam velocity can be tuned from several hundred m/s down to rest. These beams can then be loaded into static electric or magnetic traps for collision studies [77,92,94,96] or used in low velocity experiments [98,129,130].

There are two main approaches to Stark deceleration, which use distinct electrode geometries and applied voltage waveforms. The first approach, pulsed-pin Stark deceleration (PPSD), uses crossed-pin electrodes and discrete digital voltages to decelerate a pulsed molecular beam. While PPSD is relatively easy to implement using commercially available high-voltage switches to control the potentials on the electrodes, there are well-documented loss mechanisms [139,140] that reduce the efficiency of the deceleration process, particularly at low beam velocities. The second approach, traveling-wave Stark deceleration (TWSD), uses ring electrodes and sinusoidal analog voltages to create a continuously moving Stark potential well [121]. Because the molecules are always confined while being decelerated, TWSD is able to avoid many of the instabilities and inefficiencies [139] inherent in PPSD [107,118]. Additionally, by chirping the frequency of the analog voltages down to DC, the traveling potential well comes to a standstill and can serve as a trap [127]. Thus, TWSD also offers the option of efficient loading of the decelerated molecules into an electrostatic trap, resulting in greater trapped molecule densities than in previous Stark decelerator geometries.

The challenge to implementing TWSD is creating the high-voltage driving electronics. The traditional approach for creating high voltage waveforms, using transformers, fails at low frequencies, and only has a narrow tuning range. Thus, linear amplifiers must be used to drive the

traveling-wave decelerator. The necessary bandwidth, voltage, current, and heat dissipation characteristics of an amplifier are set by the particular geometry of the decelerator. Our decelerator is of similar design to previously constructed ring Stark decelerators [121,122], and has 624 tantalum ring electrodes with every eighth ring electrically connected. Therefore, TWSD requires eight or more high-voltage analog amplifiers in order to produce the chirped frequency sine-wave voltages. The rings are made of 0.04 inch diameter wire with an inner diameter of 0.16 inches. The center-to-center spacing between rings is 0.08 inches. This decelerator is designed to serve a similar role to our pulsed-pin Stark decelerator [153], and will be used with the same high-voltage power supplies, pulsed valve, and state-selective detection system used previously [106].

In TWSD, the Stark potential well must have an initial longitudinal speed that corresponds to the initial speed of the molecular beam. The speed of the potential well must then be reduced to the desired final speed of the molecular packet. Using the initial and final speeds of the molecular packet and the length of one Stark potential well, we can calculate the required bandwidth of the amplifier. The longitudinal length of one Stark potential well is determined by l , the electrode spacing, and p , the ring periodicity (i.e., the number of rings that create a single period of the sine wave). To produce a potential well with a longitudinal velocity v , the frequency of the output voltage waveform is

$$f = \frac{v}{pl}. \quad (5.1)$$

For most traveling-wave Stark decelerators constructed to date [121,125–127], including ours, l is ~ 2 mm and p is 8. For a supersonic molecular beam seeded in krypton, the initial velocity is between 400 and 450 m/s, setting the upper end of the bandwidth to 28 kHz. If the molecules are to be trapped after deceleration, the final velocity must be zero. Therefore, the high-voltage amplifier must have a bandwidth of at least 28 kHz down to DC. The amplifier presented here is designed to run from 30 kHz down to DC in order to provide flexibility with the output waveform. This bandwidth, on its own, is not extremely challenging to achieve, but there are also demanding voltage and current specifications to be met.

The peak output voltage of the amplifier and geometry of the decelerator set the depth of the Stark potential well, which in turn sets the limit on the magnitude of deceleration. The larger the amplitude of the voltage and the closer the rings are spaced, the greater the electric fields and depth of the potential well produced. The maximum achievable voltage on a Stark decelerator electrode is determined by either the high-voltage power supplies or the maximum voltage difference that can be supported between two adjacent electrodes. While a 1 mm gap can support up to 20 kV [175], any imperfections in electrode surface quality or alignment may cause arcing at lower voltages. We aim for a more conservative field of 7 kV/mm between electrodes, which requires operating the decelerator with sinusoidal voltages that have an amplitude of 10 kV. This gives a maximum on-axis electric field of 3.2 kV/mm and a Stark potential well depth of 0.57 K for deuterated ammonia molecules (mass ~ 20 amu, dipole moment ~ 1.5 Debye).

While there are methods to decelerate with lower peak voltages, such as starting with a lower velocity beam [123, 125], using multiple types of deceleration [127–130], or using a longer decelerator [126], our goal is to decelerate a supersonic beam seeded in krypton using only TWSD in a 1.25 m long decelerator. This allows us to take full advantage of the signal increase offered by the cylindrical symmetry and continuous deceleration forces without having to contend with the challenges of long decelerators [126].

In addition to the voltage requirements, the amplifier must be able to supply the current required to charge the capacitive load of the decelerator. The effective capacitive load of our decelerator is calculated to be 225 pF [95], which does not include the capacitance of any cabling between the amplifier and decelerator. We estimate our cables add around 40 pF/ft. Even if the amplifiers are mounted close to the vacuum electrical feedthroughs, ~ 2 ft of cable is still required. This adds significant capacitance to the load, which the amplifiers must drive. Using the effective capacitance of the decelerator, the maximum frequency of 30 kHz, and the maximum ± 10 kV output, we determine the amplifiers must be able to source and sink a peak current of ~ 0.5 A. In order to allow for added capacitance due to cabling and flexibility in the output waveform, the design goal for the amplifiers was set at 1.5 A.

Lastly, the design of the amplifier must be able to dissipate the heat generated by the high voltages and currents present. The electronics of the amplifiers do not prevent the amplifiers from being run continuously, however, this would create an unmanageable thermal load and require the dissipation of several kilowatts. Instead, the experiment runs at 10 Hz, and requires the amplifier to be on for ~ 6 ms in order to bring the molecules to rest. This means that an amplifier must be able to dissipate ~ 200 W when running in a pulsed mode.

Implementation of TWSD has been limited by these challenging bandwidth, voltage, and current requirements. While several commercial amplifiers are able to meet some of these criteria, none meet all. The closest is the Trek 5/80 amplifier, which has been used in several TWSD experiments [123,126,127]. It has a bandwidth of greater than 60 kHz down to DC and a maximum output voltage of ± 5 kV at peak currents of 80 mA. The custom amplifier presented here is capable of ± 10 kV peak amplitudes, peak pulsed currents of 1.5 A, and has a bandwidth of 30 kHz down to DC, which means it can produce the waveforms needed for TWSD in an apparatus that is 1.25 m in length.

5.2 Design of the high-voltage amplifier

The design of the high-voltage amplifier is based on an optically coupled, all-n-channel push-pull amplifier, shown conceptually in Figure 5.1A. Here, current is sourced to the load when the top FET (Q_T) is conducting and the bottom FET (Q_B) is non-conducting. When the opposite is true, current is sunk out of the load. Since the FETs operate out of phase, a separate optical control signal is required to drive each FET. A push-pull amplifier could also be made with both n-channel and p-channel transistors. Here, we use all n-channel devices to take advantage of their superior voltage and power handling specifications.

To determine how many n-channel transistors are required for a given application, one should consider the maximum voltage drop across a FET. If the output voltage swings the full range from the positive to negative rail voltages (i.e., $+HV$ to $-HV$), a single FET, Q_T or Q_B , will experience a voltage drop of $2|HV|$. If this is greater than what one individual FET can handle, multiple

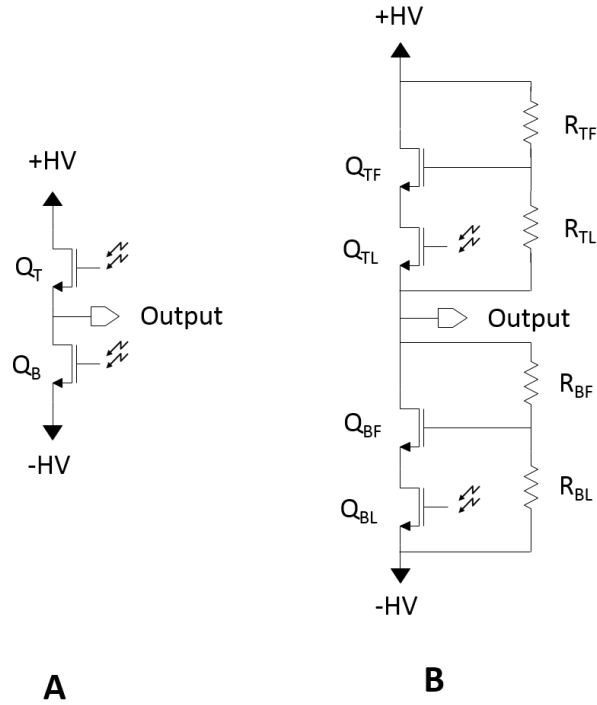


Figure 5.1: (A) An optically coupled, all N-channel push-pull amplifier. The $\pm HV$ rails and output voltage are limited to half the voltage rating of a single FET (Q_T or Q_B). (B) With multiple FETs connected in series between the HV rails and the output, a higher maximum output voltage can be achieved. The resistor network divides the voltage between the source and the output such that the voltage is shared equally across each FET.

FETs in series will be required. In our design, as multiple FETs are added in series, one FET will be driven by an external signal, while all other FETs follow in response to the leading FET. A signal-driven FET is referred to as a leader and the following FETs are referred to as the followers¹. Figure 5.1B illustrates how this is accomplished. It has a similar configuration to Figure 5.1A, with Q_{TL} and Q_{BL} as the leader FETs being driven by an external signal, but now with additional follower FETs Q_{TF} and Q_{BF} . The resistor chain on the right (R_{TF} , R_{TL} , R_{BF} , and R_{BL}) is a passive voltage divider, which allows each FET that is part of the same bank to share the voltage equally. When all FETs are off, the majority of the current flows down the resistor chain, and only a small leakage current flows down the chain of FETs. If all resistors have the same value, the

¹ In some contexts, these are referred to as master and slave circuits, but here we use the terminology leader and follower.

voltage drop across each FET is the same. If Q_{TL} begins to conduct due to a control signal applied to the gate, the voltage across it and R_{TL} will drop. This reduces the current through R_{TL} . With the current through R_{TL} less than the current through R_{TF} , the net current will flow into the gate of Q_{TF} and turn it on. Thus, whenever the voltage across Q_{TL} does not match that across Q_{TF} , a current will flow into or out of the gate of Q_{TF} to maintain the voltage sharing between Q_{TL} and Q_{TF} constituting passive, local feedback in each bank. This same process occurs if additional follower FETs are added in series.

All of the FETs in series, as well as all other components between the +HV (-HV) rail and the output, are collectively referred to as the top (bottom) bank. Each bank can be thought of as a three terminal device, with an optically controlled gate, a drain end, and source end. Within a single bank, each high-voltage FET and its associated sub circuits are referred to as one stage of the amplifier. Additional follower stages can be added to each bank at the drain end, and each follower circuit will follow the leader as outlined above.

While the push-pull design in Figure 5.1B contains the basic block structure of the amplifier, the full design of the amplifier is much more complex. As more followers are connected, additional care must be taken to ensure that all stages share the voltage equally, so that the FETs are not damaged by being over-voltaged. Failure to do this can lead to catastrophic destruction of the amplifier. The full circuits for both the leader and follower, and their relevant sub-circuits, are discussed below separately before a discussion of the amplifier control, feedback, and physical design.

5.2.1 Leader Stage

The leader stage acts as the control center of each bank. Figure 5.2 shows the full leader circuit, which consists of four primary blocks labeled with letters: a high-voltage cascode (A), global feedback driving circuitry (B), a local power supply (C), a local feedback network (D), as well as protection diodes and a few other components for biasing and current limiting. While many of the components in the leader stage float high above earth ground and experience only a few

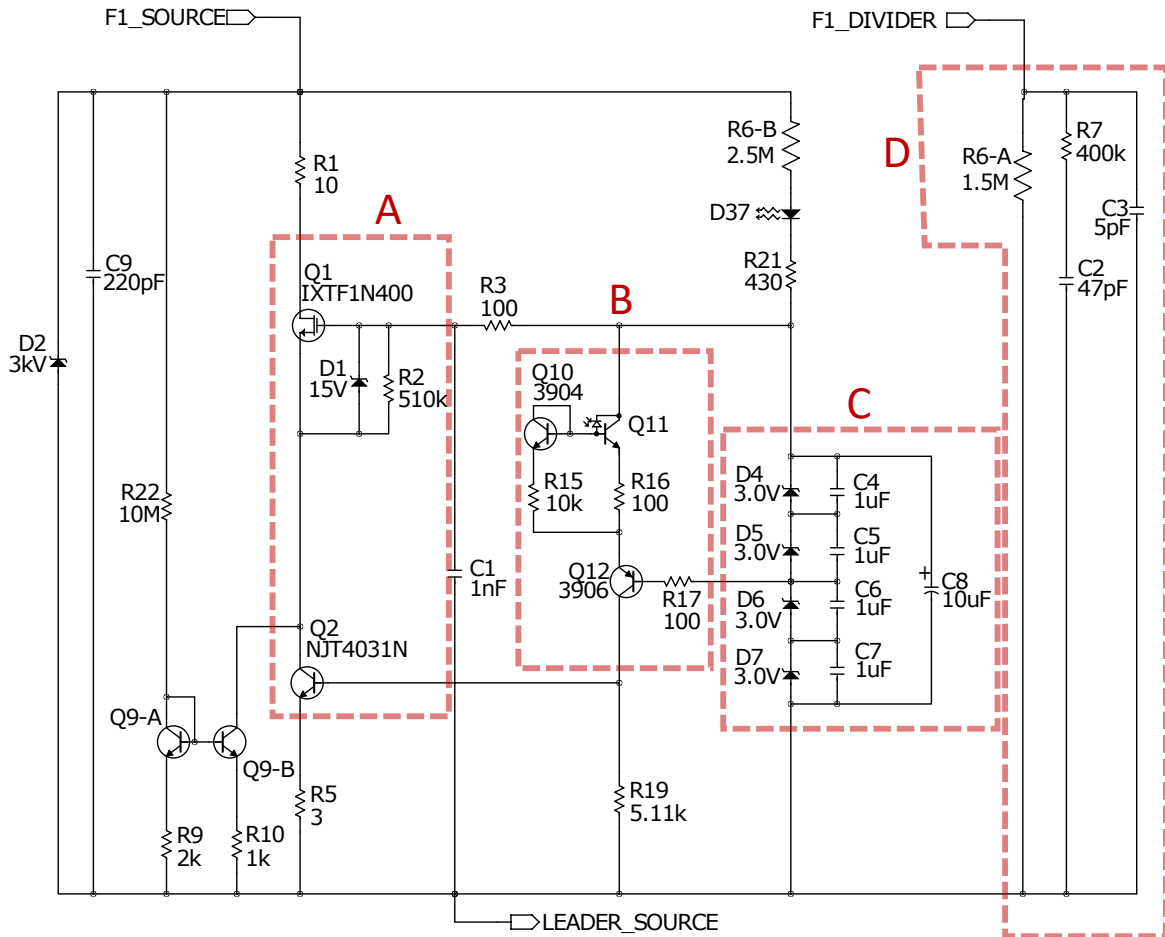


Figure 5.2: A leader stage of the high-voltage amplifier. Portions of the circuit are boxed to help identify sub-circuits, these include the A) cascode, B) global feedback driving circuitry, C) local power supply, and D) local feedback. While some components are high-voltage components, others are standard low voltage components that float up to high potentials, but do not experience a large voltage difference. Components that are high-voltage components are indicated in the text.

volts of potential difference, the HV MOSFET (Q1) of the HV cascode, the local feedback network (R6, R7, C2, C3), the protection diodes (D2), and the biasing network (R22, C9) all experience the voltage drop across the entire leader stage. These component must all be rated to withstand at least 4 kV, which is the maximum voltage rating of the MOSFETs.

A high-voltage BJT-MOSFET cascode, consisting of Q1 and Q2, is used instead of a single FET (Q_T from Figure 5.1B) to handle the primary voltage drop in the leader stage. In the cascode configuration, the BJT driven by an external signal, Q2, will swing only a diode drop between

being fully open to fully closed. This is achieved by placing a common-gate, high-voltage stage, Q1 (IXTF1400 from IXYS), between the signal-driven transistor (Q2) and the load. Overall, the cascode configuration reduces the Miller effect and increases the speed and stability of the high-voltage handling FET [176]. Additionally, the cascode reduces the amount of current required to turn on Q2, and therefore the amount of current pulled from the voltage division chain.

The global feedback driving circuitry for the main cascode is contained in Figure 5.2 Box B (Q10, Q11, Q12, R15, and R16). Q11 is a phototransistor and the end of an optoisolator. This separates the low, ≤ 1 V, control signal from the high-voltage potentials on the amplifier and any associated high-voltage noise. Q10 is a diode-connected transistor, which, together with R15, acts as a low-current pulldown for the base of Q11, enhancing the turn-off speed of the phototransistor. Additionally, Q10 and R15 provide emitter degeneration feedback, which helps to stabilize the dc bias against changes in temperature. Together, Q10-12 form an inverted cascode, which sources current to the base of Q2. The inverted cascode prevents large voltage swings at the base of Q11, reducing the amount of current needed to charge and discharge the Miller capacitance, further increasing the speed of the phototransistor. Q2 then adjusts the voltage on its collector to set the voltage drop between the gate and the source of Q1 such that Q1 and Q2 conduct the same current. When no drive current is present, R19 pulls down on the base of Q2 and closes the stage.

In order to source and sink the current required to charge the Miller capacitance of Q1, a low-impedance voltage source is required. Box C in Figure 5.2 (D4-7 and C4-8) contains a 12 V local voltage supply, which is charged through the high-voltage rail via R6b. For the range of currents drawn by Q1, capacitor backed Zener diodes are an effective low-impedance source. D37 serves as an indicator when troubleshooting to make sure the stage powers up correctly and the local supply stays charged.

The local feedback and passive voltage sharing chain is shown Figure 5.2, box D (C2, C3, R6b, and R7). It serves as a replacement for R_{TL} in Figure 5.1. The additional capacitor and resistor-capacitor in parallel help to prevent run away oscillations, while still maintaining fast feedback to the gates of the follower stages. As seen in Figure 5.1B, this portion of the circuit is relevant only

when follower stages are included.

A chain of protection diodes, D2, consisting of seven 429 V TVS diodes in series, limits the absolute maximum voltage across the leader stage to 3 kV and protects the stage from being reversed biased. The high-voltage MOSFET, Q1 has a maximum V_{DSS} of 4 kV. The protection diodes guard against any poor voltage sharing between stages, transient voltage spikes, or effects of human error during construction and testing.

The final, unboxed components provide current limiting and biasing. R5 limits the instantaneous current that flows down the cascode. Its resistance is chosen so that Q12 will go into saturation when the emitter current of Q2 is greater than 1.5 A, the maximum current for this design. Biasing of the leader circuit is important to maintain linearity and bandwidth of the optoisolator and stability of the cascode. Biasing Q1 on slightly so that it conducts a small quiescent current increases the turn-on stability, linearity, and speed. This quiescent current is created via Q9, which acts as an imbalanced current mirror. As long as there is a voltage drop across the stage, current flows through Q9a via R22. The other half of the mirror, Q9b, ensures that Q1 is slightly biased on with a low quiescent current. With rail voltages of ± 15 kV the quiescent current is 300 μA . In addition to the quiescent current, a standing current is applied just before, during, and just after the desired output signal. This is applied through the control circuitry and biases on all stages and the optoisolators to help to reduce non-linear amplifier behavior, such as crossover distortion. Typically, the standing current is set at 10 mA and is enabled 1 ms before to 1 ms after the desired waveform. Larger standing currents would further reduce nonlinearities, but would require a significant amount of power dissipation.

5.2.2 Follower Stages

The follower circuit, shown in Figure 5.3, is very similar to that of the leader circuit, and we will discuss the differences, rather than the similarities. The primary objective of the follower is to mimic the behavior of the leader stage as fast as possible. The communication between the leader and the follower stages occurs through the local feedback voltage divider network. The

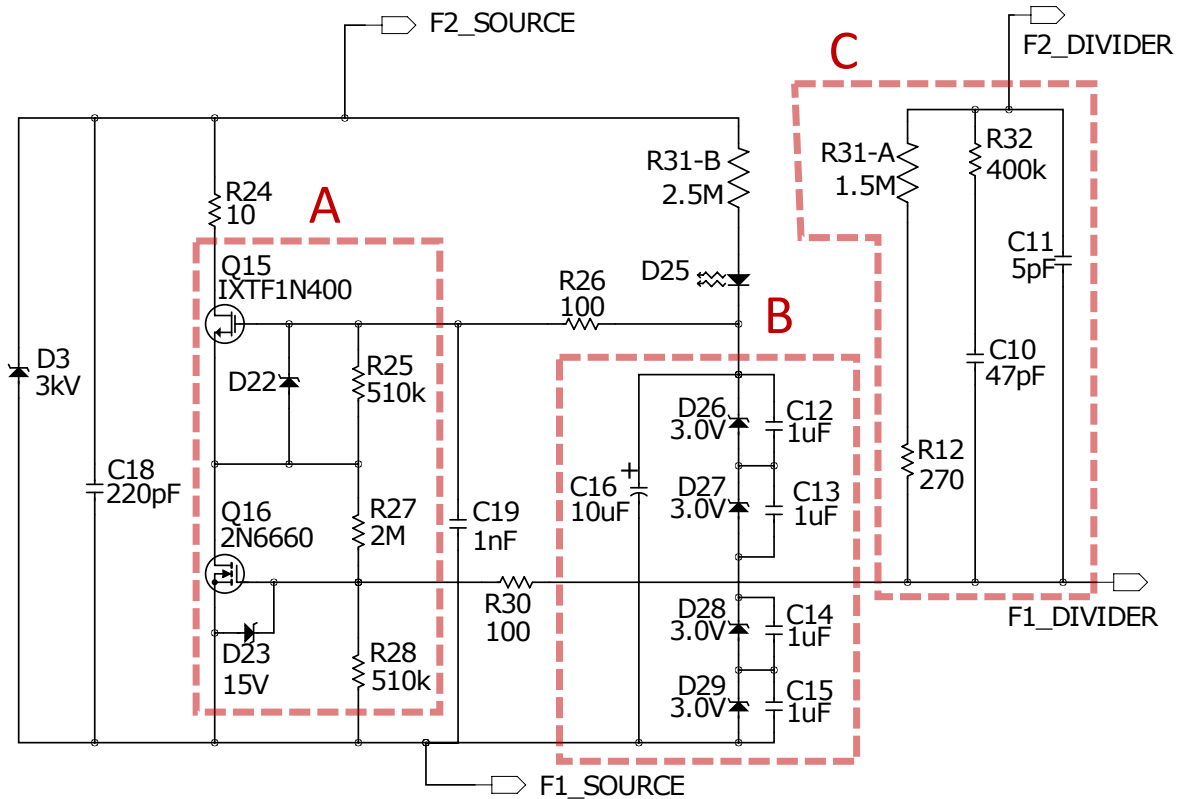


Figure 5.3: A follower stage of the high-voltage amplifier. Portions of the circuit are boxed to help identify sub-circuits, these include the A) cascode, B) local power supply, and C) local feedback.

voltage drop across each stages is determined by the current flowing through the voltage divider. Therefore, any current flowing out of the voltage divider will negatively affect the accuracy of the local feedback. However, in order for the voltage divider to act as passive feedback, it must source the current necessary to control the signal driven transistor in the high-voltage cascode, as shown in Figure 5.1. In order to reduce the current needed to turn on the signal-driven transistor, the cascode circuit on the follower (Q15-16) is a FET-FET cascode. The same IXTF1N400 high-voltage MOSFET (Q15) is used while the lower voltage, signal-driven MOSFET is a 2N6660 (Q16). Q16 then requires negligible drive current from the local feedback network. The pull-down resistor has also been changed from 5.11 k Ω to 510 k Ω , which greatly reduces the current drawn from the local feedback network to keep the gate of Q16 charged. Additionally, the FET that is driven by the external signal is faster than the BJT. For the leader stage, the linearity of the signal-driven BJT

transistor is important because any nonlinearities on the leader are corrected only through the relatively slow global feedback. For the follower stages, any nonlinearities are corrected through the fast, local feedback, so the speed of the signal-driven transistor is more important than the linearity. While global control, like the optoisolator in the leader, could also be employed on the follower circuit, only local control, via the resistor-capacitor divider chain, is used. This reduces the complexity of the follower stage and of the global feedback network overall. Therefore, there is no control signal cascode on the follower, although there is still a local power supply to drive the high-voltage FET of the cascode. Additionally, there is no biasing circuitry on the follower stages as they will follow any biasing on the leader stage.

5.2.3 Feedback and control

The amplifier contains both local feedback between the leader and follower circuits of each bank via the resistor-capacitor chain and global feedback between the external control signals and the amplifier. The amplifier is operated via a control board that uses control voltages from an arbitrary waveform generator to generate the required amplifier control signals using a PID loop for global feedback. Additionally, the control board monitors the output voltage, output current, and various other values such as the temperature of the amplifier. Ultimately, the global feedback, local feedback, and rail voltages determine the fidelity of the amplified output voltage.

In order for the feedback to function properly, the stages must be biased correctly, which is achieved by having the amplifier connected to a power supply that can keep rail voltages adequately above the desired output voltage. As long as the stages share the voltage equally, one bank with ten stages (one leader and nine followers) can support up to a 30 kV potential difference across it. Ideally, the output will not swing the full range between the rails, but if it does, one bank will see the full potential between the two rails. This means the maximum rail voltage should be half of what the transistors in series on one bank can handle. In this design, since one bank is designed to have a maximum potential difference of 30 kV, a maximum voltage on the rails is ± 15 kV.

The power supplies for the rail voltages must be backed by capacitors, so that enough current

can be supplied when the amplifiers are active. We use a Spellman SL150 and a 1 μF , 20 kV capacitor for each rail. The supplies cannot charge the backing capacitors as quickly as they are drained by the amplifiers, and as a result, the rail voltages sag a small amount during the time the output is enabled. The smallest possible capacity backing capacitor should be used in order to avoid unnecessary energy storage. In the event of an amplifier failure, all the energy may be discharged through some portion of the amplifier. The more energy that is stored, the more dangerous and destructive a discharge will be.

Under normal operation, the rails are kept a few kilovolts above the desired maximum output voltage. The amplifier operates best with at least 100 V across each stage at all times. Thus, with 10 stages on each bank, the rail voltages should be at least 1 kV above the maximum desired output voltage. In practice, especially when running in the 20-30 kHz range, having the rails 2 kV or more above the maximum desired output voltage is preferable, due to the increased current demands that will increase the voltage sag on the backing capacitors. If the rails sag below the desired output voltage, the output amplitude will be diminished and will clip until the rail voltage can recover.

A second reason to keep the rail voltages several kilovolts above the output voltage is to help produce high fidelity waveforms at higher frequencies. As more followers are added to the circuit, delays in the local feedback chain increase and affect the bandwidth of the local feedback, as well as the voltage sharing between stages. For example, a delay in the local feedback will cause a time delay between the leader and the followers. At high frequencies, this time delay may be a significant fraction of a period, affecting the final output of the amplifier. This causes the control signal to turn on/off the leader more than it would in the ideal case in order to make up the difference between the desired and actual voltage output. When the leader is driven fully open, it has only a very small voltage drop across it, and then must wait for the follower stages to catch up before the local feedback circuit can function again. While this does not lead to instabilities in the amplifier, it does lead to distortions in the output waveform.

To provide the error signal for the global feedback, a custom 20 kV 1500:1 high-voltage divider

acts as a high-voltage probe to monitor the amplifier output. The high-voltage divider is designed to have fast rise and fall times of a few hundred nanoseconds and a relatively flat frequency response between DC and 1 MHz. This feedback is enabled only when the standing current is running and the optoisolators are biased on. When not enabled, the local feedback provided by the bias chain is enough to keep the amplifier output near ground. On the control board the error signal is half-wave-rectified, summed with the DC offset that sets the standing current, and sent to each bank's optoisolator LED.

Ideally, the amplifier should have constant gain across the entire frequency range. This is mostly accomplished via careful tuning of the global feedback loop. However, if this is not possible to do to an acceptable level with a particular amplifier, a linear scaling function can be applied in the control software, which calculates the input voltage waveform in order to force the amplifier to have a constant amplitude across the frequency range. This allows a final tuning so that each amplifier will have the same output voltage regardless of variation in components and assembly.

The control circuitry also contains circuits for monitoring the output voltage and current of each bank, the cooling fan speed, and the proximity of the output voltage to the rail. If either of the last two monitors fall out of their normal operating range, the amplifier is disabled. When running multiple amplifiers, as is required in TWSD, each individual control board communicates with a master interlock. The master interlock monitors the pressure and flow of the cooling water to all amplifiers, error state of the high-voltage power supplies, and error state of all running amplifiers, all in hardware. If the master interlock detects an error on any one of these channels, it will shut down all amplifiers.

5.2.4 Physical Construction

The full amplifier consists of two banks each with one leader and nine follower circuits. Schematically, the source end of the first follower circuit connects to the top, drain end of the leader circuit at the points labeled *F1_Source* and *F1_Divider* (Figure 5.3). When connecting multiple followers, each subsequent follower, $n+1$, is connected to follower n by incrementing the

numbers at the *Source* and *Divider* points by one. Because only n-channel high-voltage MOSFETs are used, the source of the follower stage connects to the drain of the leader. Since the leader must always be at the lowest absolute voltage, it must be the source end of each bank. This results in the top bank leader being referenced to the output voltage (just like Q_T in Figure 5.1A) and the bottom bank leader being referenced to the negative high-voltage rail. The *Source* and *Divider* points of the last (furthest away from the leader) follower are electrically connected to the positive high-voltage rail (output) on the top (bottom) bank. This configuration leads to both banks being identical in construction, but a small asymmetry in operating characteristics, as current is sourced to the load by the leader of the top bank and is drained from the ninth follower of the bottom bank.

Considerations must be taken into account for the safety of the user and the amplifier including circuit layout, cleanliness, insulation, and cooling. The two banks are carefully constructed in order to minimize possible damage from capacitive coupling, corona discharge, and arcing, which all lead to degradation or destruction of the circuit. Sharp points and ragged solder joints on the circuit board lead to corona discharge, which can corrode components. In order to avoid this, the solder joints for all through-hole components are rounded off to be smooth hemispheres. The layout of the PCB uses large ground planes that are isolated from the signal traces to prevent capacitive coupling. Additionally, the components for the global feedback on the leader stage (Q10-12) are located on the underside of the PCB board to reduce crosstalk with the HV signal. Finally, the local source planes for each stage are separated by 8 mm to prevent arcing between stages.

Care must also be taken to keep the boards clean and free of dust and grease. All handling of the amplifier banks was performed while wearing gloves to prevent contamination with oils. Solder with water soluble flux was used to aid in cleaning. Once populated, the circuit boards were lightly scrubbed and placed in an ultrasonic bath to remove any flux, oils, or dust.

Even with these precautions, the amplifier can exhibit corona discharges and voltage instabilities. These can be suppressed by insulating the high-voltage circuit boards. Each bank is potted in an insulating, and thermally conductive, room-temperature-vulcanizing silicone, Insulcast *RTVS 3-95-1*. This material is chosen for its high thermal conductivity and ease of use. The silicone

seals around electronic components when it is applied, but does not adhere strongly. This allows for repairs to the board with no damage from removing the potting material. Other portions of the setup that sit at the output potential, including the output cable and feedback probe, are also potted.

The two potted amplifier banks are installed in a custom high-voltage, shielded box with active cooling (Figure 5.4). The box is constructed in two layers: an internal, insulating layer made of polycarbonate, and an external, grounded layer constructed from aluminum. Both layers of the box entirely enclose the amplifier assembly with small holes for signal pass-throughs and ventilation. The fully assembled box measures 3 inches tall, 17.75 inches wide, and 8 inches deep.




Figure 5.4: A photograph of the amplifier with the top lids removed. The cables between the rail voltage supplies and the amplifier are on the right and the output, which can connect to the decelerator, and the probe for global feedback are on the left. For the top (bottom) bank, the leader TL (BL) and last follower TF9 (BF9) are labeled.

The walls of the box where the high-voltage MOSFETs are mounted must simultaneously provide electrical isolation and cooling to the amplifier banks. The IXTF1N400 MOSFETs handle the majority of the voltage drop across each stage, and thus the majority of the power dissipation. If the amplifier runs from rail to rail at ± 15 kV and at 30 kHz, the high-voltage MOSFETs will each dissipate 340 W, which is nearly impossible to remove via water cooling. By running the amplifiers in a pulsed mode with only a few percent duty-cycle, as they would be during TWSD, the heat generated drops to roughly 10 W per MOSFET. While this thermal load is manageable, water cooling of the MOSFETs is required. Typically, the MOSFETs can be cooled directly through the isolated back tab, which is normally held at ground. However, because each stage floats many

kilovolts above earth ground, the MOSFET back tab can not be grounded. Instead, the back tab of each MOSFET is electrically connected to its source. Thus, cooling the MOSFETs requires a method that is thermally conductive, but electrically insulating, so that the source plane of each MOSFET remains isolated from all the other MOSFETs. To accomplish this, the sides of the box have boron nitride sheets set into the polycarbonate and water cooling lines are run through the grounded, metal sides. Boron nitride is a ceramic with good thermal conductivity and is also electrically insulating. The boron nitride is sandwiched between the MOSFET and the water-cooled aluminum wall, thus maintaining insulation between the large potential differences of the MOSFETs and the exterior grounded box. The boron nitride pieces are glued in place using MG Chemical's *4226 Super Corona Dope*. The interior wall and the outer aluminum wall are screwed together and thermal paste is applied between them on the boron nitride portions in order to facilitate cooling. The IXTF1N400 MOSFETs are mounted such that their back tab is in direct contact with the boron nitride.

5.3 Amplifier Performance

5.3.1 Testing Setup

The amplifier is tested on the bench using a high-voltage capacitor, C_l , connected as a test load to substitute for the Stark decelerator. A schematic diagram of the testing setup is shown in Figure 5.5. The amplifier has been tested with loads that vary from 125 pF to 500 pF in order to test different current requirements without changing the output voltage. A $10\ \Omega$ sense resistor is placed between the load capacitor and ground in order to monitor the load current. A Tektronic P6015A 1000X probe is used to measure the output voltage of the amplifier.

The amplifier is tested with single frequency and chirped sine-wave pulses 1-5 ms long, as well as with low frequency square-wave pulses. The length of the pulses is limited by the size of the backing capacitors. With larger capacitors or a power supply capable of sourcing 1.5 A, the length of the pulses would be limited only by the cooling of the amplifier. Single-frequency pulses




Figure 5.5: The test setup for measuring the output voltage and current. The 250 pF capacitor acts as the output load. The 1000X probe is used to measure the output voltage of the amplifier. The 10 Ω resistor is a current sense resistor and the 1X probe is used to monitor the output current of the amplifier. All components other than the 10 Ω sense resistor must be rated for high-voltage.

are used to characterize gain, phase-slip, and distortion across the bandwidth of the amplifier. Square-wave pulses are used to characterize rise and fall times, and time delays. Finally, while there are many types of sinusoidal pulses that can be used to decelerate molecules, linearly chirped sine-wave pulses are shown because they will be the signals used most typically with traveling-wave Stark decelerators.

5.3.2 General Specifications

The general specifications for the amplifier are listed in Table 5.1. The amplifier is designed for a peak output voltage of ± 12 kV, but has only been tested in this work to ± 10 kV in order to maintain linearity and stability. Thus, the gain is designed to be 12,000 with a maximum input of ± 1 V. The design goals for the amplifier emphasize bandwidth, output voltage, and current. The bandwidth of an amplifier can be classified in a number of ways, with the most common being the -3dB point of the amplifier's gain curve. The gain curve for our amplifier, under three different load conditions, is shown in Figure 5.6. The gain curve shows relatively little deviation from the target gain of 81.6 dB, especially for the 250 pF load, indicating that the amplifier bandwidth extends beyond the target of 30 kHz. The amplifier was not tested above 30 kHz at 20 kV_{pp} because the

local feedback may not be fast enough to ensure even voltage distribution between the stages, and a single stage may be overvoltaged. However, tests at $10 \text{ kV}_{\text{pp}}$ show the -3 dB point to be beyond 40 kHz for all load capacitances tested.

Table 5.1: Technical parameters of the amplifier.

Parameter	Value
Max input voltage range	$\pm 1 \text{ V}$
Output voltage range	$\pm 10 \text{ kV}$
Gain	12,000
Load	up to 500 pF
Peak Current	$\pm 1.5 \text{ A}$
Bandwidth	$>30^a$ or 23^b kHz
Noise	$300 \mu\text{V}/\sqrt{\text{Hz}}$
Harmonic distortion ^c	1.4 %
Supply voltage	$\pm 15 \text{ kV}$

^a Calculated from gain fall-off

^b Calculated from THD measurements. Limit 5% THD on 250 pF load

^c Measured at 10 kHz with a 500pF load




Figure 5.6: Measured gain of the amplifier. The solid horizontal line is located at 81.6 dB, which is the target gain of the amplifier. The gain was measured using an input single-frequency sine wave with a peak-to-peak amplitude of $1666 \text{ mV}_{\text{pp}}$.

While the gain curve indicates a bandwidth greater than 30 kHz, the quality of the amplifier

output begins to deteriorate beyond 18 kHz. One measure of output signal fidelity for high-voltage amplifiers is Total Harmonic Distortion (THD) [177]. THD compares the amplitude of the Fourier component of the peak frequency against the amplitude of all higher-order harmonics, and indicates the distortion from an ideal, single-frequency wave. If V_n labels the n th harmonic amplitude, then THD is defined as:

$$\text{THD} = \frac{1}{V_1} \sqrt{\sum_{n=2}^{\infty} V_n^2}. \quad (5.2)$$

Distortion of the sine wave will appear “smooth” if the amount of power in the higher harmonics decreases with increasing harmonic number. An example of two output waves with their Fourier transforms is shown in Figure 5.7. Although the 30 kHz waveform shows a significant THD of 11%, the distortions appear smooth due to the decreasing power of the higher-order harmonics. An analysis of the THD across the frequency range shows that the THD stays below 2% for frequencies below 18 kHz. The THD then rises sharply for frequencies above 18 kHz to a final value of 11% at 30 kHz. For Stark deceleration, the primary requirement is that the amplifier reach the required voltage and operate at the required frequency. Smooth distortions will slightly change the edges of the Stark potential well, but are not likely to strongly affect the deceleration of the molecules, particularly since these distortions occur when the molecules are moving at high speeds, where the molecular packet is less sensitive to non-ideal potentials.

Because of the smooth nature of the distortion and the primary reliance of Stark deceleration on the absolute amplitude of the signal, we will consider both the THD and the gain when determining the bandwidth of the amplifier. For applications where amplitude is a primary concern, the bandwidth is greater than 30 kHz. For applications where waveform fidelity is of concern, a bandwidth of 18-20 kHz is more relevant. The relatively large amplitudes of the even-ordered harmonics indicate an asymmetric distortion between the positive and negative sides of the wave, which is evident in the time-domain waveform. This is due to the asymmetric nature of the all-n-channel-MOSFET amplifier construction noted earlier.

The output voltage of the amplifier was tested at ± 10 kV. The gain curve shown in Figure




Figure 5.7: Measured output sine waves and their respective Fourier transforms for 4 kHz (a, c) and 30 kHz (b, d) input sine waves on the 250 pF load capacitor with ± 13 kV rails. The x-axes are scaled such that there are the same number of periods (3) and harmonics (15) shown for each of the frequencies. (a) The time domain and (c) frequency domain response for a 4 kHz input sine wave have a calculated THD of 0.7 %. (b) and (d) are the same for a 30 kHz input sine wave, with a calculated THD of 11.0 %.

5.6 was taken with a constant amplitude input sine-wave of $1666 \text{ mV}_{\text{pp}}$, and shows a near-constant gain of 81.6 dB, giving a constant output of $20 \text{ kV}_{\text{pp}}$ across all frequencies. As tested, the amplifier reached a peak current of 860 mA, while driving the 500 pF load at 30 kHz. This is significantly above the 500 mA needed to charge the 250 pF load of our Stark decelerator.

Additional tests were performed with a square wave to measure rise and fall times and time delays. The results of these tests are summarized in Table 5.2. The values for rise and fall times highlight the asymmetry introduced through the all-n-channel-MOSFET construction of the amplifier banks. While the fall time, controlled by the bottom bank sinking current from the load, is dependent only on the capacitance of the load, the rise time is dependent on both the capacitance and on the output voltage. This indicates that the output current that the top bank can source is dependent on how far the output voltage is from the rail voltage.

Table 5.2: Small-signal amplifier parameters. Rise and fall times, t_r and t_f , and delay d were taken with a 1 kHz square wave with an output of V on a load capacitance of C_l . The rails were held fixed at ± 11 kV.

V kV _{pp}	C_l pF	t_r μ s	t_f μ s	d μ s	S^a V/ μ s
10	125	3.9	3.4	2.8	2590
12	125	4.8	3.4	2.8	2530
10	250	4.2	4.4	3.2	2320
12	250	5.3	4.4	3.2	2260
10	500	4.9	6.2	3.9	1630
12	500	5.6	6.2	3.9	1920

^a Slew rate is calculated from peak amplitude and the larger of t_r or t_f .

5.3.3 Voltage Sharing




Figure 5.8: Ratio of voltage dropped between the final follower stage (F9) and the leader stage (L) on the bottom bank for three different frequencies. The voltage sharing for high frequencies is poor at low voltages, but stabilizes as the output voltage increases. All data were taken with rail voltages of ± 6 kV. The inset shows the waveforms for the voltage across the leader and final follower at 30 kHz and at 8 kV_{pp}. The point at which these data were taken is highlighted on the main plot with the red star.

Voltage sharing between the stages of the amplifier via the local feedback is critical to the stability of the amplifier in order to assure that no single stage is over-voltaged under any conceivable operating conditions. In order to assess the quality of voltage sharing between stages of one bank, the voltage across each stage is measured by separately probing the voltage at the source end and then the drain end of each stage and then subtracting the signals during post-processing. The source and drain ends of the stage need to be measured separately, or crosstalk between the measurement probes' leads can inject current into the stage and drastically alter the stage's functionality.

The largest voltage difference occurs between the stage that is closest to the rails and the

stage that is closest to the output. For the bottom bank, these are the leader and ninth follower stages, respectively, and vice-versa for the top bank. The behavior of the two banks is similar, and the following discussion refers only to the bottom bank. The ratio between the voltage across the ninth follower stage to the voltage across the leader stage is shown in Figure 5.8. Data were taken at 4, 20, and 30 kHz, and at amplitudes from 2 kV_{pp} to 8 kV_{pp}. Higher amplitudes were not tested because measuring voltages across stages requires probing inside the amplifier while it is running, which requires the potting material to be removed from the circuitry. This, in turn, decreases the stability of the amplifier, and increase the likelihood of corona effects at high voltage.

Figure 5.8 shows that while higher frequency signals cause the amplifier to share poorly at low voltages, increasing the output voltage increases the stability of the local feedback network at high frequencies. The voltage across the final follower stage reached about 1.33 times the voltage across the leader. The inset shows a sine fit of the measured voltage on both stages. While the amplitude difference is clear, there is very little phase shift between the two signals.

The voltage sharing ratio is a measure of how well the voltage is distributed across the entire amplifier bank. If the ratio is too large, the ninth follower stage could experience too large a voltage drop, and could become overvoltage. To determine if the measured sharing ratio of 1.33 is low enough to allow the amplifier to produce a ± 10 kV sine wave with ± 13 kV rail voltages without overvoltage any stage, we consider the following calculation.

The mean voltage across a stage, \bar{V} , is defined as one tenth of the voltage across an entire bank, and is the ideal case for the amplifier. A voltage sharing ratio different from unity indicates that the voltage across the leader and ninth follower are offset from \bar{V} . Measurements of this offset, V_{offset} have shown that it is symmetric about \bar{V} for the leader and the ninth follower. Thus, the voltage drops across the leader and ninth follower are $V_L = \bar{V} - V_{\text{offset}}$ and $V_{F9} = \bar{V} + V_{\text{offset}}$. The sharing ratio is then defined as

$$R = \frac{V_{F9}}{V_L} = \frac{\bar{V} + V_{\text{offset}}}{\bar{V} - V_{\text{offset}}}. \quad (5.3)$$

The maximum value for the voltage across a single stage, V_{\max} , is set by the protection diode chain to be 3 kV. The ninth follower stage is at a higher voltage than the leader stage, and may potentially rise above the diode threshold voltage. This sets a bound on the maximum allowable voltage offset of

$$V_{\text{offset}} < V_{\max} - \bar{V}. \quad (5.4)$$

Combining Equations 5.3 and 5.4 gives an upper bound on the acceptable value of the voltage sharing ratio of

$$R < \left(2 \frac{\bar{V}}{V_{\max}} - 1 \right)^{-1}. \quad (5.5)$$

For the amplifiers to safely produce an ± 10 kV sine wave with rail voltages of ± 13 kV, each bank must be able to withstand a maximum of 23 kV, which sets \bar{V} to 2.3 kV. R must then be less than 1.875 in order for no stage to be overvoltaged during normal operating conditions. The measured R value of 1.33 is well below this maximum value. Therefore, we do not anticipate this small voltage sharing error to limit the performance of the amplifier.

5.3.4 Chirped sine waves




Figure 5.9: A linearly chirped voltage waveform from 30 kHz to DC at ± 10 kV. Fewer periods than would be used for TWSD are shown for clarity, and three start-up periods are appended to the front of the wave packet in order to prevent large start-up current spikes.

Ultimately, the amplifier will be used for TWSD and must be able to output a chirped sine-wave voltage with uniform output amplitude across the frequency range. A sine-wave chirped

from 30 kHz to DC at ± 10 kV output voltage is shown in Figure 5.9. Here, for clarity, a more aggressive chirp is shown than would typically be used in TWSD. A typical deceleration waveform in our decelerator would have 78 periods in 6 ms. A DC hold is also demonstrated, as this would be required for trapping the molecules. Additionally, a three period start-up sequence, where the amplitude is increased linearly to the operating voltage, is appended to the front of the waveform to prevent the large current spikes necessary to start a sine wave at arbitrary phase from ground. The high frequency part of the waveform shows distortions similar to those in Figure 5.7, while the lower frequency part is free of distortions.

5.4 Conclusion

A new type of Stark deceleration technique called traveling-wave Stark deceleration promises to greatly increase the number and density of cold, slow molecules produced compared to traditional pulsed-pin deceleration. The challenge to implementing TWSD lies in creating high-voltage analog amplifiers with demanding specifications. The bandwidth, voltage, and current requirements needed to bring a supersonic beam seeded in krypton down to rest in 1.25 m are not available in any commercial high-voltage amplifier. The amplifier presented here is capable of output voltages up to ± 10 kV, instantaneous currents up to 1.5 A, and a frequency range from 30 kHz down to DC, which will allow for deceleration of a molecular beam with an initial velocity of >450 m/s down to rest. We have quantitatively measured the voltage sharing error between stages and the total-harmonic distortion. We demonstrated that despite small deviation from ideal performance, the amplifier can operate over the full design frequency range. Additionally, the flexibility of these amplifiers can allow for not only traveling-wave Stark deceleration, but also efficient trap loading and manipulation.

5.5 Individual amplifiers

Each individual amplifier setup consists of an amplifier box (which contains the two banks), a high-voltage probe, a breakout box, and a control board. The circuits for all these components

can be found in Appendix A. Figure 5.10 shows the amplifier, probe, and breakout box. The 20 kV 1500:1 probe is separate from the amplifier box and plugs into the amplifier via an Amphenol high voltage connector. The probe can be attached via screws so that it cannot be accidentally removed. The amplifier is run by a control board, of which the front panel is shown in Figure 5.11. Two control boards are mounted in each panel mounted box, and the left edge of the second control box can be seen in the Figure. The control box converts the input commands and feedback signals into drive signals for each amplifier bank. The inputs are a manual enable switch, TTL enable input, control voltage (labeled VIN), and +12 V power for the amplifier fans. The two enable signals dictate when the global feedback for the amplifier is active. The three monitors provide low voltage output monitors for the output voltage (VOUT) and top and bottom bank current (ITOP and IBOT). Lastly, the control board also contains knobs for adjusting the proportional and integral gain for the amplifier global feedback and eight indicator LEDs. From top to bottom the error LEDs (red) indicate: fan error, fan error latched, rail error, and rail error latched. The errors latch so that any failures that are too fast to observe by eye can still be caught. Once the amplifier detects an error it turns off until the error is no longer present and the clear button has been pressed. Error monitors can be disabled via jumpers on the control board. The four status LEDs (green) indicate: amplifier “ok”, amplifier enabled, standing current enabled, and tuning mode. Each goes green if that state is true. The amplifier ok state indicates that there are no detected errors.

A breakout box communicates between the control board and the amplifier/probe. It reads in the voltage from the high-voltage probe, contains the probe compensation circuit, and separates out the control and monitor signals signals from a 50-pin cable into three 8 pin connectors. Ethernet cables are used for the 8 pin connections because they are a convenient shielded cable and commercially available in many styles. The three ethernet cables contain the control and monitor signals for the top bank, bottom bank, and fans.

Each assembled amplifier must be tested and tuned on the bench before it is placed on the decelerator. A picture of the testing setup diagrammed in Figure 5.5 is shown in Figure 5.12. For




Figure 5.10: A potted high-voltage amplifier with its probe (left) and unattached breakout box (top). Holes have been cut into the potting material to allow for placements of probes for current tuning. The probe (left side) plugs into the amplifier via an Amphanol high voltage connector. When in operation, the probe is connected to the breakout box using a BNC cable (not shown). On the right side of the amplifier box, two Amphanol high voltage connectors provide connections to the positive and negative high voltage rails and two fans provide additional cooling. Three ethernet cables (not shown) make the connections between the the breakout box and the amplifier. The ethernet cables provide control signals for the top and bottom bank and carry power power to the cooling fans.



Figure 5.11: The front panel of the amplifier control box. The manual enable switch, TTL enable SMA input, and V_{in} SMA input are noted by red arrows. The manual enable switch and TTL enable SMA input together control the output state of the amplifier. The desired chirped sinewave control signal is set via the V_{in} SMA input.

bench testing, in addition to the amplifier, probe, breakout box and control box from Figure 5.10, a low-voltage power supply, a high-voltage power supply, backing capacitors, arbitrary waveform generator (AWG), and a digital delay generator (DDG) are required. The low-voltage supply provides the +12 V fan power, and the high-voltage supply and backing capacitors provide the rail voltages for the amplifier. The AWG creates the sine-wave control voltages, which are fed to the control board. The DDG triggers the AWG and creates the TTL amplifier enable signal. The enable signal must begin ahead of the AWG trigger so that the standing current begins ahead




Figure 5.12: The high-voltage amplifier must be first tested on the bench to ensure equal sharing between stages and faithful input/output voltages. For bench testing, in addition to the amplifier, probe, breakout box and control box from Figure 5.10, a low-voltage power supply, a high-voltage power supply, backing capacitors, AWG, and a DDG are required. Since often the amplifier top is open during testing so that probes can be placed, the setup should remain as clean as possible.

of the sine-wave signal. An additional fan (not shown) plugged into the same power strip as the high-voltage supplies is often helpful as a audible reminder that the high-voltage is on. The power strip should be switched off whenever the high-voltage supplies are turned off.




Figure 5.13: The amplifier can be carefully probed at both test point and at jumper connections. Here two probes have been placed in the jumper connections in follower 9 of the top bank. This is not an ideal configuration since the probes are relatively close to one another (see text).

A special lid for the amplifier was created for individual testing so that probes could be inserted into the amplifier (Figure 5.13). This setup allows access to the interior of the amplifier so extreme care must be taken when testing. Testing should never be completed alone. The test bench

should remain as clean as possible since the amplifier lid is often open. Each amplifier stage can be probed in test points (labeled TPx) and at jumper connections (labeled Jx) where x represents the component number which is different on each stage. To probe test points, a Tektronic P6015A 1000X probe is connected to a high-voltage clip cable, which can be attached to the test point. To probe a jumper connection, the Tektronic probe is connected to a stiff high voltage wire that matches the gauge (19 SWG) of the jumper connector. Since the wire that runs between the probe and board is not insulated, cross-talk can occur and change the measured signal. Thus it is important to keep the high voltage wires from crossing and away from other cables. Cross talk can lead to incorrect measurements, and it is important to remember that probing the amplifier circuit can easily change its behavior. While multiple probes can be placed on the amplifier as shown in Figure 5.12, care must be taken as to where these probes are placed because what is measured may not be what is actually happening due to the probe impedance changing the normal behavior of the amplifier. For a full step by step procedures for tuning and testing see sections 6.5.2.1 and 6.5.3 of [95] respectively.

5.6 Insulation

5.6.1 Amplifier bank

The original amplifier design presented in [95] is prone to current instabilities at output voltages larger than ± 6 kV. In addition to the oscillation caused by the voltage limiters discussed in section 5.8, two other types of oscillations show up. The first kind is persistent and occurs on the falling edge of the current oscillation. It was found on all sixteen amplifiers that have been built. The second type of oscillation causes a dramatic dip in the output current 10-20 μ s wide, but only occurs sporadically.

The first persistent current oscillation is dependent only on the output voltage and not on the voltage across a stage (number of enabled stages) or current (load capacitance). For example it will appear with ± 8 kV output voltage using ten stages on each bank but not at a ± 4 kV output

voltage using five stages on each bank regardless of the output load. It is frequently observed on the falling edge of the current oscillation which corresponds to the positive peak output voltage of the amplifier, as shown in Figure 5.14a. This oscillation is most easily seen in pictures of the oscilloscope since infinite persistence traces are very difficult to save to data files. Infinite persistence is used to capture all the different current oscillation behaviors and to show that the oscillation does not occur in the same location in time on each shot. It is also useful when debugging the amplifier, since it will catch spurious behavior that occurs only in a small percentage of the iterations of the output.



(a) Current oscillations issue



(b) With the addition of SF₆ the current oscillation issue is almost entirely suppressed.

Figure 5.14: ± 8 kV output without (left) and with (right) the amplifier box filled with SF₆. The output voltage is shown in green and the output current as measured across the sense resistor is shown in blue.

Figure 5.15 shows a close-up on the type of oscillations seen in Figure 5.14a. This type of oscillation had a very repeatable frequency character and examinations of multiple oscillations at different amplitudes all resulted in similar frequency responses. The oscillations begins with a fast and quickly decaying 130-140 MHz oscillation follow by a 12 MHz slower decay.

In each shot of the output, only a few oscillations occur. The picture in Figure 5.14a contains 97 shots of the experiment and many oscillations are overlaid. Although the oscillations do not always appear in the same place in time on each shot, they remain contained to the positive falling edge of the current. The largest oscillations always occur on the first one or two periods of the output current sine wave. This was true even at different output frequencies. There is also a




Figure 5.15: A close-up of the persistent oscillations seen in Figure 5.14a. The oscillations have two characteristic frequencies, an initial fast response at 130-140 MHz and a slow response at 12 MHz.

smaller oscillation that appears on the negative rising edge, but since it does not grow significantly with the output voltage, we were not very concerned about it. These oscillations typically started at a threshold output voltage of ± 6 -8 kV and their behavior had some hysteresis. As the output voltage was slowly increased there would not be oscillations below a threshold voltage. Once the oscillations began, if the output voltage was adjusted to be below the previously observed threshold the oscillations would remain. However, if the amplifier was turned off for a few minutes and then turned on again at a voltage below the threshold there would be no oscillations.

Although both of these types of oscillations were not previously reported in the version of the amplifier in [95], these instabilities occur frequently on all sixteen of the amplifiers that have been built. These oscillations do not appear to affect the output voltage at moderate output voltages (± 6 -8 kV), but since they grow in amplitude at a faster rate than the output current amplitude we have chosen to address them before they caused significant problems. It is suspected that these current instabilities are caused by corona discharge. In a very dark room, a few faint blue corona spots could be seen particularly at the corners of the vertical FET wall.

Originally, the amplifier circuit boards were coated with *Super Corona Dope* (MG chemicals), a one-part insulating coating. However, this did not improve the corona discharge situation. It also made repairing the amplifier boards more difficult. A soldering iron can burn through the clear and hard cured *Super Corona Dope*, but this had to be done in a very well ventilated area. In discussions with the manufacture, we learned that although it has a rating of 4100 V/mil, this is not linear and does not build with thickness. The coating was also designed for coils and not PCB boards so it was not suited for our applications. We also tried *Silicon Conformal Coating 422B* (MG Chemicals) to address the corona discharge, but again found no improvements.

While in the initial design it was suitable for the circuit board to just be in air inside the specially designed box, an insulating gas or potting material significantly improved the performance of the amplifier and greatly reduced the current oscillations. Two dielectric mediums that are often used in high-voltage circuit breakers and switches are oil and SF₆ gas. We did not test oil due to the difficulty of containing it and making repairs to anything that has been submerged in oil. SF₆



Figure 5.16: 50 shots at a 1 Hz repetition rate of ± 12 kV output at 5 kHz using an amplifier filled with SF₆. The output voltage is shown in green and the output current as measured across the sense resistor is shown in blue.

has a much higher dielectric strength than either air or dry nitrogen. The amplifiers were tested in SF₆ and it dramatically reduced the oscillations. It did not make the oscillations disappear entirely, but this may have been due to the difficulty of keeping the amplifier boxes gas tight.

Figure 5.14b shows an amplifier with the same output voltage and current as Figure 5.14a, but with the addition of SF₆. For this ± 8 kV output, the current oscillation issue is almost entirely suppressed and indicates that corona discharge and other high-voltage insulation and isolation issues are the cause of this type of current oscillation. Thus with improved insulation, the current oscillation should be suppressed. With SF₆ the amplifier was able to go up to our design specification of ± 12 kV without significant current oscillation issues (Figure 5.16).

SF₆ was tested in two configurations, initially by placing the amplifier in a sealed glove bag and later by sealing the box shut with tape and adding a tube to evacuate the inside of the amplifier and fill it with SF₆. Neither of these methods produce very gas tight seals, but the SF₆ remains in the amplifier box long enough (a few minutes) to make measurements. The tape method produces a weaker seal, but it was a easy modification, used less SF₆, and allowed for better access. In the glove bag configuration, in order to limit the number of wires and tubes that entered and exited the bag, the entire breakout box, all ethernet connectors, and the probe were also placed inside. If the setup was left for several hours, oscillations would usually return, indicating that SF₆ had leaked out and had been replaced by air. All tests with SF₆ were run with the entire test setup placed

Table 5.3: RTVS 3-95-1 properties

Viscosity	10,000 cps
Pot life 25°C	1.5 h
Cure time 25°C	24 h
Share A hardness	65
Thermal conductivity	1.25 W/(m K)
Dielectric strength	500 V/mil

in a fume hood.

While SF₆ showed promising results for solving the oscillation issues, it is not a practical long-term solution. The amplifier box design requires two high-voltage inputs, two high-voltage outputs, and two ethernet connections, which carry the low-voltage control signal, and one more ethernet to power the fans and monitor their status. Due to the difficulty of constructing a gas-tight box that can accommodate these connections, in addition to the cooling requirements, we chose to use a potting material to encapsulate the amplifier banks instead of SF₆ gas. The requirements for selecting a potting material were high dielectric strength, high thermal conductivity, and low viscosity. We selected Insulcast RTVS 3-95-1, which has the bonus feature that the silicone compound does not bond with most materials and can be removed for repairs without damaging the electrical components. This means potting is not permanent and allows for repairs, which are tedious but possible. When removing the potting material, care must be taken to not rip up any components (most typically the test points). Although the material does not bind to any of the components, it does seep under anything that is not soldered down. Thus, if it is removed in large chunks, the components can be ripped off the board. The only material that we found that RTVS 3-95-1 did not easily release from was glass. It also does not cure properly when it comes into contact with various adhesives. The properties of mixed RTVS 3-95-1 are listed in Table 5.3.

Initially, only a thin coating of RTVS 3-95-1 was used on each amplifier bank (Figure 5.17). However, with just a thin layer on the banks the oscillations did not diminish. Placing the thinly potted banks back into a sealed box, and adding SF₆ improved the behavior of the current indicating that there was not sufficient insulation. In the end, the banks were potted into bricks (Figure 5.18).





Figure 5.17: A thin layer of RTVS 3-95-1 potting material on one amplifier bank. This was not enough potting material to stop the oscillations and adding SF_6 gas to the box improved the insulation



(a) Potted amplifier bank immediately after it is released from the mold.



(b) Potted amplifier bank with high voltage connection points cut out.

Figure 5.18: Two fully potted amplifier banks. Cut out high-voltage connection points can be seen on the front-left corner of the picture on the right.

This has the added benefit of stabilizing the optoisolators, which have very fragile solder connections due to the very fine (45 gauge) wires, and which must support to ethernet communication receptacle. At this point, adding SF_6 to the amplifier box no longer had any affect. Although the fully potted amplifier did not perform quite as well as the amplifier in SF_6 , it was still a significant improvement over a non-potted amplifier.

RTVS 3-95-1 is a two part room-temperature-vulcanizing silicone. Part A (red) and part B

(clear) are mixed in a 20:1 ratio. A new can of Part A always arrives slightly separated and must be well mixed (at least 10-15 minutes) before weighting out the two parts. Part B does not require any premixing. When first mixed, the material has the consistency of thin honey and takes 24 hours to cure, though a bit longer to reach the final stiffness. Mixing introduces many air bubbles, so the mixture is evacuated in a bell jar. this produces a smooth finished product. The evacuated mixture must be poured into the mold slowly to ensure that no air bubbles are trapped in the underside of the circuit board. Once released from the mold, the finished potted amplifier bank is very smooth. Any wrinkles or soft spots indicate large air bubbles underneath, these can be repaired by cutting open the air pocket and filling it again with RTVS 3-95-1.

The amplifier boards are potted while attached to the outside wall via 'L' shaped brackets and nylon screws. A mold is constructed of PVC and polycarbonate pieces and held together with modeling clay and double sided tape. This allows the mold to be easily released, wiped down, and used again. Although the tape causes the corners of the mold not to cure, any soft corners or edges can easily be shaved off with a razor blade once the mold is removed. Small plugs were placed in the high-voltage rail and output connection points before pouring in the potting material in order to keep the connectors clean from the silicon. After the silicon cures, the plugs are cut out and removed (Figure 5.18b). This can be done fairly easily with a small razor blade.

Potting the amplifier banks helped to eliminate many of the first type of oscillation but a second type of sporadic oscillation was still present. The second type of oscillation was found to occur when the pink high-voltage wire that comes with the Amphenol Alden high-voltage connector leans against polycarbonate. This insulated wire (pink) is 22 AWG and rated up to 30 kVDC. However, the connector (white) is only rated to 25 kVDC. While the first type of oscillation was a high frequency and symmetric oscillation, this one has a high frequency spike followed by a much slower dip, Figure 5.20. The dip always followed the direction the current was already going, but with a steeper slope. Whenever this oscillation occurred, a simple rearranging of the high-voltage wires inside the amplifier box or output cable could resolve this issue. Keeping the length of pink high voltage wire inside the amplifier relatively short reduces these oscillations because there is less

wire to accidentally curl up and touch the interior polycarbonate box.

This oscillation with a dip also occurs when pressure is placed on the output cable that connects the amplifier and the load. For the input and output cable on the amplifier, we use in-house built cables made from Amphenol Alden A401QX and A800QX connectors, the QX wire that comes with the Amphenol connectors or other 30 kVDC HV wire, additional insulation, and a ground shield. Figure 5.19 shows a cross section of the cables used with the TWSD setup. The high-voltage wire is placed inside 3 mm wall Tygon tubing. The cable is then surrounded by ground braid, which shields against noise and pickup and provides protection should all insulators fail. The left cable in Figure 5.19 has been used for high-voltage switching in our PPSD experiment with good success. However, with the sine-wave voltages from the amplifier output, there are occasional current and voltage dips. These can usually be eliminated by adjusting the angle of the output cable, which suggest that the issue again stems from the position of the interior high-voltage wire. To remedy this, the space between the Tygon tubing and wire is filled with the same RTVS 3-95-1 potting material as the amplifier circuit board (right cable Figure 5.19). The input cables to the amplifiers that just sit at the high-voltage rails are not potted.

The increased insulation adds extra capacitance of 20-40 pF/ft to the cable. The amplifiers are mounted such that they are as close as possible to the vacuum electrical feedthroughs, but are still mounted practically. Each cable from the amplifier to the decelerator is ~ 2 ft. From Figure 5.6, even with the additional capacitance and current draw, this will not significantly affect the behavior of the amplifier.

The high-voltage probes are also potted, which further reduces current noise. While some probes were potted with RTVS 3-95-1, we found it easier to pot them with Hysol ES1001 since it is easier to prepare and has a lower viscosity. The high-voltage probes are very reliable once built and all tuning is done externally, thus potting them with a non-removeable material is not an issue.




Figure 5.19: Cross section of the custom high voltage cable used in the TWSD setup. A RTVS filled cable (right) is used if it carries an AC voltage, non-potted cables (left) are used to carry DC voltage, such as from the Spellman voltage supplies to the amplifier.



Figure 5.20: Current (blue) dips caused by proximity to polycarbonate. Usually current oscillations do not affect the voltage (green) on the decelerator, but when the current issues are severe enough they also cause droops and bumps in the output voltage.

5.6.2 Amplifier walls

High voltage isolation is not only a concern for the amplifier circuit board, but also for the amplifier box. The two layer box consists of an interior primarily polycarbonate box for electrical isolation and an exterior grounded aluminum box. Two sides of the polycarbonate box not only had to provide voltage isolation, but also cooling for the IXTF1N400 FETs. A piece of boron nitride, a thermally conductive, but electrically insulating ceramic, is set into two sides of the polycarbonate wall to allow for heat exchange between the FETs and the exterior aluminum walls, which are water cooled. Although Super Corona Dope was found to be not useful for the amplifier boards, it was used to seal the space between polycarbonate and boron nitride in the FET walls of the box. It serves as a dielectric glue that holds the boron nitride in place. Figure 5.21a shows a picture of just

the polycarbonate and boron nitride wall. It is then affixed to the aluminum cooling wall by screws through the aluminum wall and into blind tapped holes in the polycarbonate. These screws should not be any more than finger tight, tightening them will flex the polycarbonate wall and cause cracks in the Super Corona Dope seal. Thermal paste is applied between the walls to facilitate thermal conductivity. Building these walls is a tedious process that involves filling the gap between the boron nitride and polycarbonate in many steps. The boron nitride and polycarbonate have a step cut into it in order to increase the path length between the FET wall and the grounded exterior wall. Due to the small space that needed to be filled, the surface tension of the Super Corona Dope causes many air bubble to form, thus it must be applied in thin layers and the bubbled popped each time before it cures. Each layer must begin to cure (~ 6 hours) before the next layer can be applied. This is most easily done with a 23-25 gauge hypodermic needle so that the bubbles can be released as each layer is applied. While this method creates wall that are able to sustain a 20 kV difference for at least 5 minutes, there are probably filling materials for this application which could be easier to use. If the walls were to be remade, the electrically insulating but thermally conductive panel could possibly be made out of the same RTVS potting material used to insulate the amplifier instead of boron-nitride. However, this would make them incredibly fragile and other more rigid potting materials should be explored.



Figure 5.21: The interior insulating polycarbonate wall with boron nitride inserts. Each side wall uses two boron nitride pieces (A). The black scorch mark on the right-hand side shown a point where the wall failed. A close up of the inset boron nitride (B) shows the step that increases the path length between the aluminum wall and the back plane of the FET.

Once assembled, the walls must be tested to ensure that they can withstand 20 kV. It is easiest to test the walls without thermal paste initially. This way the walls are much easier and less messy to repair. Once the walls have passed the isolation test, thermal paste should be applied and the wall tested one last time. The test set up should use no backing capacitors and have a low current limit. The outside wall MUST be grounded at all times. The voltage should be increased slowly so that if the wall fails, it fails at the lowest voltage possible.

Wall failures are audible and visible if the lights are low. Once the wall is no longer at high voltage, the failure point should be visible when the polycarbonate wall is removed from the aluminum wall. The failure point will usually leave a fairly noticeable scorch mark on the step portion of the boron nitride and a small mark on the aluminum wall. Occasionally the Super Corona Dope will turn yellow and brittle at the point of failure. Wall failures are almost always caused by air bubbles in the Super Corona Dope. While from the surface there may not appear to be air bubbles near the failure point, once the top layer of Super Corona Dope has been removed, there almost always will be an air bubble underneath. Wall failures can be repaired by carefully removing the Super Corona Dope in the vicinity of the failure and refilling the opening.

5.6.3 Optoisolator Failures

Even with the potting material, the amplifiers still encountered some unfortunate high-voltage issues. In all but one of these events, the negative high voltage rail shorted to the ground plane on the ethernet board via the current monitoring optoisolator D21/Q52. This would usually annihilate Q2, Q13, and Q14 and the surrounding resistors. Figure 5.22 shows the result of one of these failures.

Each optoisolator consists of a transmitting LED and a receiving phototransistor. Although the OPI120 optoisolator comes in a package, each of these subcomponents is listed separately on the circuit diagram. Q51/Q11 is the optoisolator for driving the leader circuit and D21/Q52 is the optoisolator for monitoring the amplifier current. Although the component Q51 is a LED, it is listed as a 'Q' transistor component because of a historical error. We keep it as Q51 so that it is compatible with [95] and because the D51 is already a different component (one of the TVS diodes




Figure 5.22: A side on (left) and top down (right) view of an amplifier that had a short to ground. Potting material has been removed for visual access to the damaged portions of the board. The side on view shows how the board was bent in the explosive event. The top down view shows several of the scorched components.

in the protection chain).

These failures always destroyed D21/Q52 and occasionally Q51/Q11. Since it appeared that D21 was the culprit in these failures, the bottom bank current monitoring circuit was removed from the circuit by shorting the planes at TP4 to TP7. This is most easily accomplished by connecting the -HV rail (TP7 and TP8) connection points to the MASTER_SOURCE plane (TP4) via the pads originally designated for the brass shield (J20), Figure 5.23. This results in D21 conducting no current and both leads constantly sitting at the negative rail voltage. Even with this modification, D21/Q52 would occasionally fail at rail voltage greater than ± 10 kV, even though it was within the specifications of the optoisolator, which is rated to ± 15 kV and possibly higher if encapsulated. One of these failures occurred on the test bench and not on the decelerator. This was less destructive due to the smaller backing capacitor ($1 \mu\text{F}$ instead of $5 \mu\text{F}$) used on the bench and thus less energy was dissipated through optoisolator and circuit board. In this case, only D21/Q52 failed and caused minor damage to Q51/Q11 and no damage to the board. After replacing both optoisolators the amplifier returned to normal operation. Although previous amplifier failures were violent and destructive, this suggests that D21/Q52 was the source of the trouble and all other damage was collateral.

The OPI120 has an input-output isolation voltage rating of ± 15 kVDC and if suitably encapsulated it can be increased to ± 25 kVDC. While the top bank optoisolator must isolate between




Figure 5.23: The current monitoring optoisolator is shorted out by connecting the -HV rail (TP7 and TP8) connection points to the MASTER_SOURCE plane (TP4) via the pads originally designated for the brass shield (J20).

ground and the oscillating output voltage, the bottom bank optoisolator always isolates between ground and the negative high-voltage rail. The bottom bank optoisolator is always the one that fails even though it is the one sitting at DC, while the top bank optoisolator oscillates. The next iteration of these amplifiers will replace the OPI120 with OPI150, which has an input-output isolation voltage rating of ± 50 kVDC. This will require some modifications to the driving cascode circuit due to the lower gain of the OPI150 and physical modifications to the amplifier mounting system to account for the increased height of the OPI150.

One amplifier failure did not result in the a short to the ethernet ground plane and was barely audible. Instead, it was only marked by a rail error on the control board. When disassembled, all $1/2$ W 10Ω resistors on the drain end of the big FETs on the bottom bank were destroyed. It seems that this failure traveled up the chain of follower stages instead of to the ethernet board ground plane. This failure mode was not investigated in depth since it occurred only once.

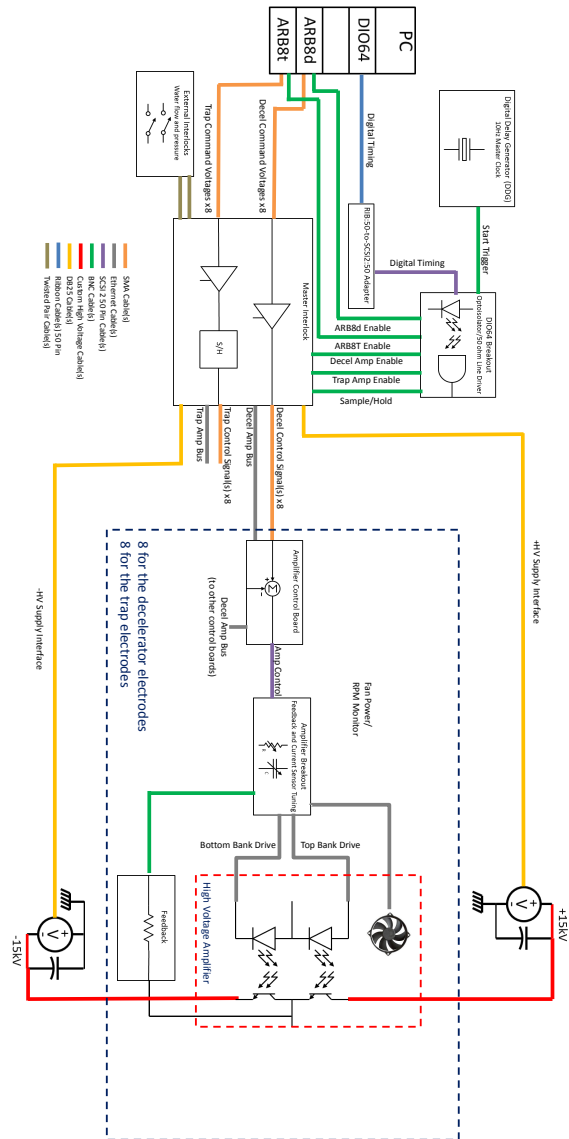


Figure 5.24: A simplifier schematic of the amplifier electronic and control connections. Modified from Figure 6.43 in [95]. One amplifier contains everything inside the red-dashed rectangle. Everything required to run one amplifier is contained inside the blue-dashed rectangle.

5.7 Operating multiple amplifiers

In order to run the full decelerator, eight amplifiers, each with their own control, feedback probe, and breakout box, must be connected properly to the decelerator and computer control system. The full control system for running TWSD is shown in Figure 5.24. Components inside the amplifier box itself are contained in the red dashed rectangle. Inside the blue dashed rectangle is everything required to run one amplifier. This selection of parts must be repeated for the number of amplifiers being used. Only one of all parts outside of the blue dashed rectangle are required. Working approximately left to right, a computer generates the analog and digital signals to run the decelerator. The digital signals control when each component of the experiment turns on and the analog signals control the amplifier waveforms. A digital delay generator is the master clock for the experiment and controls the frequency at which the decelerator runs. A master interlock box checks for errors and only allows the amplifiers to output high voltage only if everything else in the experiment is running properly. Lastly, two ± 15 kV Spellman SL (SL 2000) power supplies and are backed by two $5 \mu\text{F}$ backing capacitors provide the rail voltages to all amplifiers.

5.7.1 Digital and analog signals

Control of the experiment is accomplished via a PC running LabVIEW and with two *GaGe 8152* (eight channel arbitrary waveform generators with ± 1 V outputs) cards and a *Viewpoint Systems DIO64* (64 channel FPGA-based digital I/O) card. The computer calculates two sets of timing sequences, one analog and one digital, based on the user inputs for the desired experiment. The analog signals are produced by the GaGe boards and control the amplifiers. The GaGe boards produce up to a ± 1 V output. With a gain of 12,000 on the amplifiers, this is sufficient to drive any desired output on the amplifiers up to ± 12 kV. The digital sequence sets the timings for the molecular beam valve, enabling and starting the amplifier waveforms determined by the analog boards, detection lasers, and TOFMS voltage switches. Similar to the PPSD experiment, the digital control outputs are optically isolated and buffered with 50Ω line drivers to help reduce

interference from high-voltage noise and to prevent damage to the computer should there be a high voltage failure.

The full amplifier control system is designed to support 16 analog amplifier control signals produced by two eight channel GaGe arbitrary waveform generators. Eight of these channels are dedicated to the amplifiers attached to the decelerator and eight are dedicated to the amplifiers for the trap. Since we perform only deceleration and not trapping in the experiments in this dissertation, only one of the GaGe boards is used. Figure 5.24 includes all connections required for trapping, however if trap amplifiers are not going to be used, these portions of the setup can be ignored. While theoretically any eight rings of the decelerator can be used as a trap for a short amount of time, if trap times longer than ~ 10 ms are required, it is best to use rings that are electrically isolated from the decelerator so that the entire decelerator does not need to be held at the trap potential during trapping. Using separate amplifiers also means that there is more flexibility when trap loading since as the molecule approach the trap, there are actually 16 ring electrodes that can be at 16 different potentials, which increases the trap loading possibilities.

The computer calculates the full deceleration waveform (Equation 2.14) required for a given initial and final velocity. However, this full sequence may not always be used. In order to detect molecules ahead of the timing molecule, the laser must fire before the full deceleration sequence is completed. The decelerator cannot be on when the laser fires because the resulting ions will be steered by the strong fields from the electrodes at the end of the decelerator. The electric field of the TOFMS plate voltages, which steers the ionized molecules to the MCP detector, will be modified by the decelerator. To avoid this, the deceleration sequence is abbreviated and the decelerator grounded right before the detection laser fires. The full calculated waveform is referred to as the deceleration sequence, while the actual, possibly abbreviated sequence, is referred to as the GaGe output.

The entire experiment is timed off of a digital delay generator (DDG), which provides the master clock. The DDG triggers the DIO board, which sets when to enable the amplifiers and when to start the GaGe boards. The majority of the TWSD experiments are run at 5 Hz. Although the

experiment was designed to run at 10 Hz, running at 5 Hz reduces the duty cycle on the amplifiers and increases the time for the backing capacitor bank to fully recharge. If the capacitor bank does not fully recharge on each cycle, the rail voltage will not be maintained. The amplifier output will then be distorted, and the master interlock may kick in and turn off all amplifiers.

Figure 5.25 shows the digital timings to run a decelerator experiment as described in section 2.5. Here, all times indicated with a upper case T are times referenced to the beginning of the timing sequence, while times indicated with a lower case t are referenced to another trigger. T_0 represents the time when the DIO64 board is triggered by the DDG. Two triggers are required for the Nd:YAG that pumps our dye laser, one for the flash lamps and one for the q-switch. There is a set delay between the flash lamp trigger and the q-switch trigger, t_q . The q-switch trigger occurs at T_1 and all other timings are referenced from it. The main time that will be varied in a deceleration experiment is the detect time of the molecules, t_{detect} . The valve is fired at a time $T_1 - t_{detect}$ and the rest of the experiment is timed off the valve. For optimal deceleration, molecules loaded into the decelerator should have an initial velocity that matches the initial velocity of the chirped sine wave. This means the decelerator should turn on after a time $t = \frac{d}{v_i}$ where d is the distance from the valve to the decelerator entrance and v_i is the initial mean velocity of the molecular packet. Although this time can be calculated, it will not be accurate since there is a delay between the time the valve is triggered and when the valve actually opens, and as a result, this delay is determined experimentally. This delay is called t_{GaGe} since it is the time in which the GaGe board is triggered after the valve.

Although the GaGe output is not a digital signal, it is shown in Figure 5.25 since it is used to calculate the amplifier enable signal. The GaGe output is a sine wave signal that controls the amplifier output. Only one GaGe channel is shown here but in practice there is one GaGe signal per amplifier. The GaGe output is triggered on the rising edge of the GaGe enable signal. It continues until either the end of the deceleration sequence or until T_1 , whichever comes first. In Figure 5.25, we have shown what would be the entire deceleration sequence, but since it ends after T_1 , the sequence is truncated. The portion of the sequence shown with the dashed line would be at

ground. If t_{detect} were increased to a point where the entire deceleration sequence occurred before T_1 , then the GaGe output would go to ground after it finished its sequence.

While all other digital signals trigger systems that begin on a rising edge, making the exact pulse length not crucial, the length of the amplifier TTL enable signal is very important. This signal sets when the amplifier has its global feedback and standing current on. It must turn on before the desired sine-wave output and remain on until after the sine-wave output ends in order to ensure that the amplifier maintains linearity. Typically, the amplifier is enabled 1 ms before the desired output and is kept on for an additional 1 ms after the output has finished. The length of the amplifier enable signal is determined by the actual GaGe output signal and not by the full calculated deceleration sequence. Therefore, if the GaGe signal is prematurely cutoff due to a short t_{detect} , the amplifier enable signal ends 1 ms after T_1 (as shown in Figure 5.25). If on the other hand, the GaGe output ends before T_1 , the amplifier enable signal ends 1 ms after the GaGe output ends.

Picking the time for T_1 is not critical except that it must be sufficiently after T_0 to ensure that the amplifier enable signal does not need to occur before T_0 ,

$$T_1 > 1\text{ms} - t_{GaGe} + t_{detect_{max}}. \quad (5.6)$$

Here $t_{detect_{max}}$ is the maximum detect time that will be used in the experiment. All triggers are referenced from T_1 , the time that the q-switch fires because the pulsed YAG laser must run at 10 Hz. If all timings were referenced from the time that the valve fires or when the decelerator turns on, the laser frequency would change each time the detect time for the experiment changes. In order to keep the laser as stable as possible, it is best to have all timings referenced to the laser flash lamp or q-switch triggers.

For the case of running the experiment at 5 Hz, the laser must still fire at 10 Hz. In Figure 5.25, two extra triggers are added to the flash lamp and q-switch triggers in order to trigger the laser a second time during the 5 Hz experimental cycle. The triggers are placed 100 ms after the initial triggers to maintain the 10 Hz laser repetition rate. There is a ~ 30 ns jitter in the DIO64

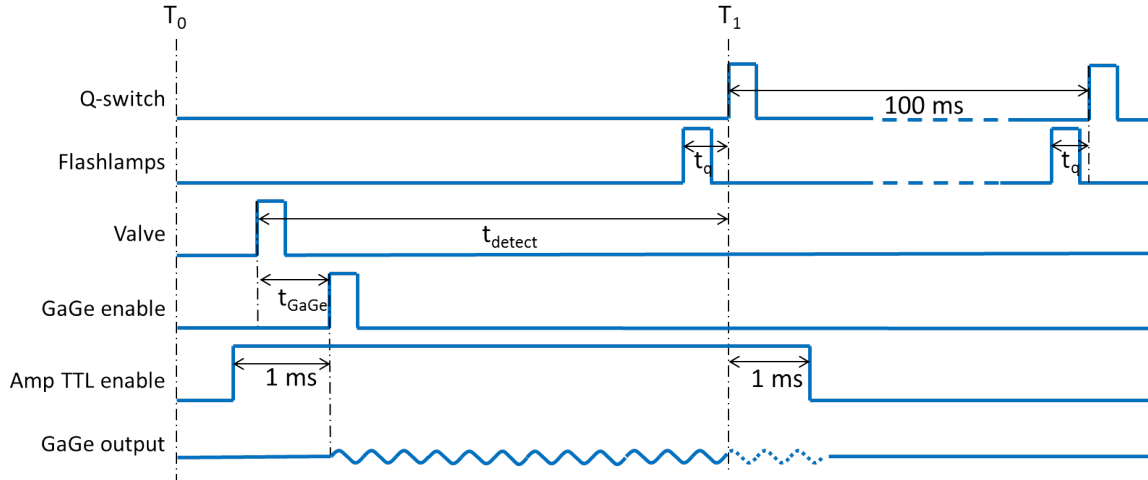


Figure 5.25: A schematic of the digital timings. T is the time referenced to the DDG and t is relative time between triggers. Although it is an analog and not digital sequence, one GaGe output is shown at the bottom. The timing of the GaGe sequence is dependent on the digital sequence.

board between receiving the start trigger and beginning its sequence, thus it is best to have the laser q-switch and flash lamp triggers on the DIO64 board instead of the DDG to eliminate the jitter from the experiment. However, if pressed for channels on the DIO64 board, the triggers for the laser can be run off of the DDG. In this case, the DDG is run at 10 Hz and the DIO sequence is extended to be longer than 100 ms even if the experimental sequence is significantly shorter than 100 ms. By extending the timing sequence, the DIO board will be triggered only once every other time the DDG sends a trigger.

5.7.2 Master Interlock

The signals from the GaGe boards are passed to the master interlock box, which serves as the global control of whether the amplifiers are enabled. Under normal operation, the master interlock passes all the analog signals from the computer to each amplifier control board and monitors for any errors. Once it detects an error, all signals to the amplifier control boards are halted until the error is addressed. The master interlock monitors the status of each amplifier, external conditions such as the rail voltage and water pressure, and offers access to additional utility functions of the power supplies.

While each individual amplifier monitors its own rail and fan error, the master interlock checks each control board via a daisy chain bus communication scheme. Each amplifier has a unique binary address (0000-1111). The master interlock cycles through the range of amplifier addresses using a 555 timer (U25) and a sequential 4-bit counter (U24). Each control board reads the instantaneous address and if it matches that of the control board, then that particular board communicates any errors to the master interlock. If any amplifier has an error, the master interlock will ground all analog control signals and display which amplifier has an error via four LEDs on the front panel. Like much of the amplifier design, the bus communication is carried on an ethernet line. In addition to the four pins required for the address and one for the error bus, the other three pins of the ethernet daisy chain from the master interlock to the control boards contain the enable signal, +HV_RAIL and -HV_RAIL signals. The last two signals communicate the Spellman voltages to the amplifier.

The master interlock contains a sample-and-hold ability to sample the desired end trap voltage and hold for a time longer than the GaGe board can output. The GaGe boards do not have enough memory to generate both a fast chirped sine wave (30-0 kHz in ~ 6 ms) and then hold a constant value for seconds as required for trapping. The sample and hold functionality is performed by an SMP-04 4-channel sample-and-hold IC on channels 8-15 on the master interlock board. The sample-and-hold function can be enable with digital signal from the DIO board. This function was not used for the experiments described in this dissertation since no trapping was performed. With no sample-and-hold control signal into the master interlock, the eight trap channels can be used exactly like deceleration channels. Although the master interlock is designed to have amplifiers 0-7 run as the decelerator amplifiers and amplifier 8-15 run as the trap amplifiers, it is possible to swap the amplifiers around. It is permissible to plug an amplifier with an address between 0 and 7 into the trap bus or to use the trap channels to control decelerator rods. The master interlock does not check the decelerator/trap bus against the amplifier addresses.

In addition to the amplifier error signals, the master interlock monitors the flow and pressure of the cooling water. It also communicates with the two Spellman voltage supplies and interlocks the

voltage supplies to the water cooling. If the cooling water does not have adequate flow or pressure due to a clog, leak, or building issue, the Spellman supplies will be shut off. This indirectly grounds the amplifier input signals since the sagging rails voltages will cause a rail error on each individual control board. Turning off the Spellmans to turn off the amplifiers is not a suggested method of turning off the amplifier except in unusual or emergency circumstances.

The master interlock also includes a few options for controlling the Spellman supplies including disabling front panel voltage and current control (so that the Spellmans can be controlled by a separate analog signal) and the ability to link the two supplies together so that the -HV supply always follows the +HV supply. Although this is useful for ensuring that the output of the amplifier stays at zero when the supplies are first turned on, it is not recommended to use this function of the master interlock box. We found that due to the long cables required to connect the master interlock and Spellman supplies in our current setup, slaving the Spellmans together leads to a voltage sag on the -HV supply and is not recommended unless the 25-pin cable (yellow in Figure 5.24) can be greatly shortened.

5.7.3 Enabling amplifiers

In order to output a voltage, each individual amplifier must be enabled and have a control signal from the GaGe board. The enable signal turns on the global feedback, and provides the amplifier with the standing current explained in Section 5.2. To enable an amplifier, the manual enable switch on the control box must be in the on position and the amplifier must receive a TTL enable signal (Figure 5.11). This TTL enable signal can come from the front panel enable SMA input (typically used when testing individual amplifier) or through the ethernet daisy chain (typically used when running amplifiers on the decelerator). Only when both of these signals are high does the amplifier fully enable and begin amplifying the control signal. The enable LED (second from the top in the green LED stack) will be on when both signals are high. Under normal operation when the amplifier is enabled for 5-8 ms, the blinking is easily seen by eye.

While testing the amplifier, the easiest and fastest way of turning the amplifier on and off

is via the manual enable switch while keeping the TTL enable signal and control signal active. While this is an acceptable method of testing one amplifier it is not the optimal method of running multiple amplifiers attached to the decelerator. Instead, when running the decelerator, all amplifiers should be manually enabled before any control signal is sent. This ensures that global feedback is enabled any time voltage can be applied to the decelerator. Even if only one rod of the decelerator has a voltage on it, all amplifiers on all other rods should have their feedback enabled and output forced to ground. This prevents any decelerator rods from being at an undefined potential. For example if AMP1 is applying voltage to rod 1 and AMP2, connected to rod 2, does not have its feedback enabled, rod 2 will capacitively couple to rod 1 and will float up to an arbitrary potential between ground and the voltage on rod 1. However, if the feedback on AMP2 is enabled and the control signal grounded, AMP2 will actively try to keep rod 2 at ground. Since there is a limited bandwidth to the feedback, rod 2 has residual oscillations when its control signal asks it to be at ground.

Figure 5.26 shows the voltage on a rod when the control signal for the amplifier is at ground, while the amplifier has the feedback both on and off, and an adjacent rod is oscillating at ± 1 kV. The red traces shows when feedback is never enabled, and the black trace shows when the feedback is enabled starting at 0. The drift away from 0 V when the amplifier feedback is off is due to small amplifier stage sharing error. When global feedback on an amplifier is not enabled, voltage sharing across the two banks is predominantly dictated by the resistor divider chain and there may be small differences between stages since the chain is created by ten 1% resistors. Without the amplifier feedback enabled, the rod, which is supposedly at ground, swings ~ 150 V and is centered near 0 V. With the amplifier feedback on, the voltage on the rod swings by ~ 100 V centered around 0 V.

In order to limit the voltage swings on the rods when some amplifier are on and others are off (such as when first turning on the decelerator), a software turn-on procedure was implemented. Instead of using the manual enable switch to control which amplifiers are outputting a non-zero voltage, a software control is used. The amplifier output should be controlled via the GaGe board. The recommended procedure is:

- (1) Calculate desired digital control signals and TTL enable timings
- (2) Update DIO board with calculated TTL sequence
- (3) Set all GaGe outputs to ground (via LabVIEW)
- (4) Turn on Spellman high voltage supplies
- (5) Manually enable each amplifier using the manual enable switch and make sure that the output voltage remains near 0 V using the V_{out} monitor
- (6) On the computer, change the GaGe board output to desired control signal

This procedure was started because there were concerns that current from the decelerator rod could be pushed onto the bottom bank of disabled amplifiers. While all the followers maybe able to share any excess current or voltage via passive feedback, the leader would not. The base of the control BJT in the leader is not connected to the feedback chain so the leader turns on only when it is given a control signal. Since it does not have an active control signal, the leader cannot adjust the voltage across it and could conceivably be over voltaged. This was considered to be one potential cause for the bottom bank optoisolator failures discussed in section 5.6.3. While this no longer appears to be a possible cause, this procedure is maintained since it is an overall safer way to run the amplifier and decelerator.

Lastly, on the decelerator, each amplifier is turned on with a three period ramp up sine wave (Figure 5.9). The turn on waveform is prepended to the deceleration waveform such that the amplifier turns on smoothly from 0 V. This eliminates any current spikes that will occur if a capacitive load is forced to an arbitrary voltage, as would be the case if each rod was forced to the required voltage (from ground) at the moment the timing molecule enters the decelerator.




Figure 5.26: The voltage on a decelerator rod, measured with a high-voltage probe, when the amplifier control signal is grounded, but the adjacent amplifier has an oscillating voltage. The adjacent rod has a voltage oscillation of ± 1 kV. On and off refer to whether the manual enable switch, and thus the feedback is on or off. In the case of the feedback being on (black), when it is enabled at $t=0$ the output becomes centered near 0 V and remains there until the feedback turns off. Additionally having the feedback on reduces the amplitude of the oscillation of the “grounded” rod. Offsets from 0 V when the amplifier is not enabled occur because of small deviations in the resistor chain.

5.8 Amplifier Changes

Several changes have been made to the amplifier design since [95]. Here, we detail why various components have been added, changed, or removed. Updated circuit diagrams for the complete amplifier circuit and a concise list of changes can be found in Appendix A.

5.8.1 Component value changes

The Tables 5.4, 5.6, and 5.5 give the changes to specific component values on the amplifier board, control board, and breakout board respectively. These components have remained in place, and only their value has changed.

On the amplifier board (Table 5.4) two resistors have been replaced, R21 (on the leader right above the local power supply) and R159 on the small optocoupler board. This change to R21 only needs to be made if the voltage limiter on the leader is to be used. For reasons detailed below, we no longer use the voltage limiters (Q7 and Q8). If the voltage limiter is used, R21 sets when

Table 5.4: Changes to amplifier boards resistor values. The change to R21 is optional if Q8 or Q7 are not installed, see the text for an explanation.

Component number	Original value	Updated value
R21	470 Ω	430 Ω
R159	300 Ω	100 Ω

the voltage limiter begins to conduct. With Q8 and Q7, removed the value of R21 becomes rather insignificant since it is in series with half of R6, which is 2.5 M Ω . On the original circuit diagrams, the part number for the protection diodes on big and little FETs (D1, D12, D22, D23, D40...) are not listed, we use MMSZ4701.

On the amplifier breakout board (Table 5.5), C4 and C6 are part of the high-voltage probe compensation circuit. This circuit is individually tuned for each probe following the procedure outlined in section 6.5.1 of [95]. Although this is done primarily via the trimcaps C5 and C7, 470 pF and 100 pF are more useful places to start for C4 and C6 respectively. If these values still do not produce a faithful square wave, the next two most likely changes are to increase C4 by adding an additional 68 pF or 100 pF capacitor in parallel or to reduce the value of C6+C7 by removing C6.

On the master interlock the 10 k Ω resistor before the ADG613 switch (R34, R56, R59...) was removed and replace with a 0 Ω resistor. Additionally pin 3 of the op-amp (U22, U60, U64...) was tied to ground using a 1 M Ω resistor.

On the amplifier control board (Table 5.6), R13 was changed to have the correct input impedance from the master interlock box, which has a high impedance output. On the other hand, if running without the master interlock box and directly from the GaGe boards then a 50 Ω input impedance is required. However, the master interlock should only be bypassed when testing on the

Table 5.5: Changes to amplifier breakout boards

Component number	Original value	Updated value
C4	330 pF	470 pF
C6	180 pF	100 pF

Table 5.6: Changes to control boards

Component number	Original value	Updated value
R13	51 Ω	1M Ω
R35	200 k Ω	390 k Ω
R38	200 k Ω	390 k Ω
R166	100k Ω	10k Ω
R167	10k Ω	7.5k Ω
R168	5.11k Ω	7.5k Ω
R170	10k Ω	7.5k Ω
R171	100k Ω	10k Ω
R173	5.11k Ω	7.5k Ω

bench and not when the amplifiers are connected to the decelerator. R35 and R38 are part of the optocoupler LED drive circuit and help set the standing current. They were adjusted so that we could set 10 mA as the typical standing current. The remaining changes on the control board are to resistors that set the rail proximity limit and must be adjusted to fit the power supply being used.

5.8.1.1 Rail Proximity Monitor

The rail proximity limit sets how close the amplifier output voltage can get to the rail voltages before the amplifier errors. The circuit for this section of the control board is given in Figure 5.27. R166, R167, R168, R170, R171, and R173 can be adjusted to adapt for various power supplies output voltage monitors and desired rail proximity limits. The Spellman supplies that we use in our lab have an output monitor voltage of 0-10 V where 10 V is the maximum output of a particular supply. Thus on a 15 kV supply, 5 V on the output monitor would mean the supply is at 7.5 kV, but if it were a 20 kV supply, 5 V on the output monitor would mean the supply is at 10 kV.

On the control board, U5 and U41 are set up as differential amplifiers that determine how far the output is away from the rail voltage. In order to function properly, the correct scaling between the output monitor and the rail monitor is required. Taking the positive high-voltage rail differential amplifier as an example, U5 and its surrounding resistors, R165-168 must have the correct ratio of resistances to provide the proper scaling of $+HV_RAIL$, $OUTPUT_MONITOR$, and

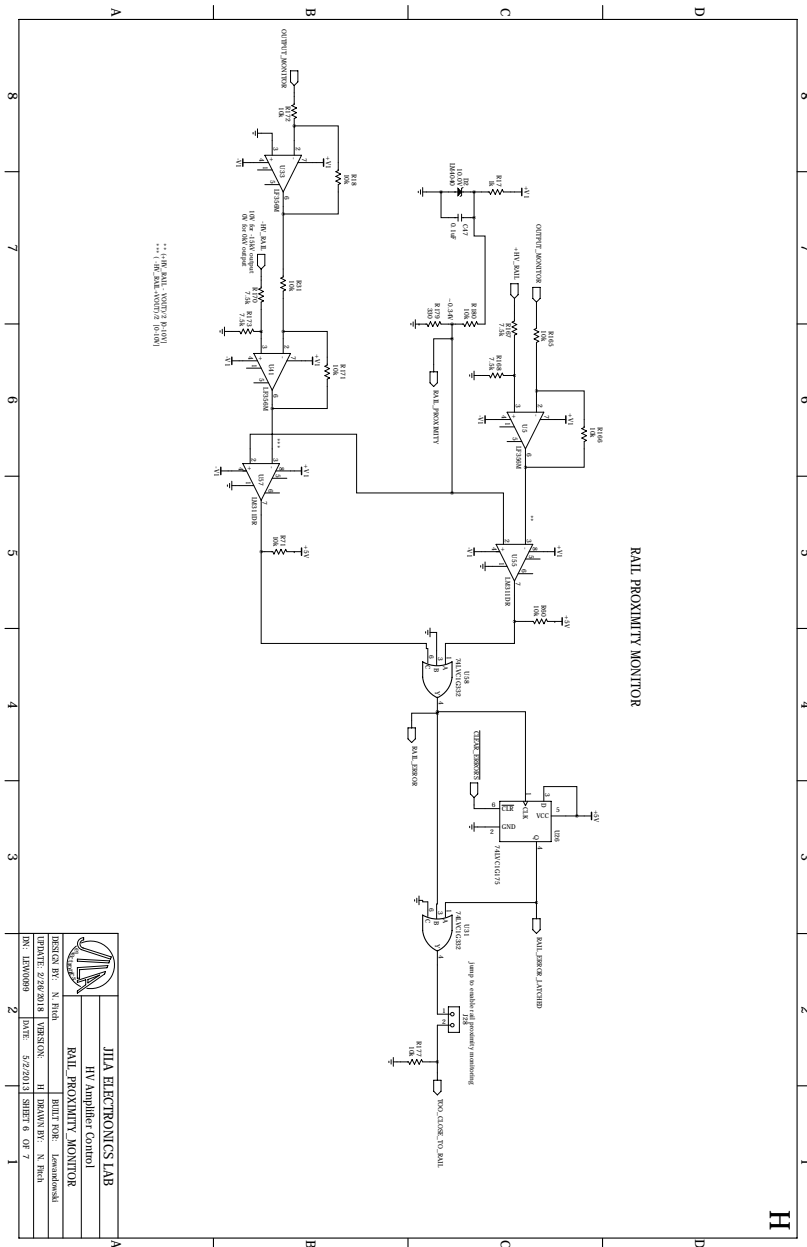


Figure 5.27: Control board rail monitor circuit.

the output. The input at R167 is the output monitor from the Spellman supply, which is scaled by 1 V/1.5 kV. The input at R165 is the output monitor voltage from amplifier. The output voltage of the differential amplifier is given by

$$V_{out} = \frac{(R166 + R165) R168}{(R168 + R167) R165} V_{OUTPUT_MONITOR} - \frac{R166}{R165} V_{HV_RAIL}. \quad (5.7)$$

Since a ± 5 kV supply has the same scaling as our output monitor (1 V/1.5 kV) on the control board, both coefficients in front of $V_{OUTPUT_MONITOR}$ and V_{HV_RAIL} should be one. This results in $R165 = R166$ and $R167 = R168$. We set $R165 = R166 = 10 \text{ k}\Omega$ and $R167 = R168 = 7.5 \text{ k}\Omega$. The rail monitor is set by the voltage at $RAIL_PROXIMITY$. We set the error limit at 500 V. Although ideally the output will remain at least 1 kV and preferably 2 kV away from the rail, this gives some buffer for noise pickup on the control board, which can sporadically cause the rail error to go off. Thus the rail monitor errors only when the amplifier is forced to go significantly above its operational limits. The scaling of the $OUTPUT_MONITOR$ voltage is 1 V/ 1.5 kV, which means a rail limit at 500 V corresponds to 0.34 V at $RAIL_PROXIMITY$.

If 20 kV Spellman supplies are used instead, the coefficient in front of V_{HV_RAIL} should be 4/3. This means R168 needs to be twice as large. So when working with 20 kV Spellmans R168 an R173 must be changed to 20 k Ω .

5.8.2 Parts not in use anymore

The voltage limiters, bottom bank current monitors, and the brass shields on both banks were eliminated from the original amplifier design. The voltage limiters were removed because the NPN transistors were easily damaged under regular operation. Once damaged, they would cause a small, fast current oscillation at the exact same location on every period of the current sine wave. It is unclear what was causing the damage to the transistor. While the oscillations remained small, this component was removed due to its propensity to break. Since each stage contains one, or two in the case of the leader, transistor as part of the voltage limiter, this makes 22 possible failure points on each amplifier.

The bottom bank current monitors were disabled because they were the most common point of failure while running the amplifiers (see Section 5.6.3). Since the optoisolator also provides support for the ethernet connector, the component was not removed but instead shorted out. This is accomplished by connecting the plane at TP4 to TP7, see Figure 5.23. In order to prevent this issue, one option is to replace the OPI120 with OPI150 optoisolators, which have an input-to-output isolation voltage of ± 50 kVDC. However, this would not be just a single component swap. Due to the difference in current transfer ratio, the driving portion of the leader circuit will need to be modified. This modification is currently being investigated.

The brass shield (connected at J20), which was originally placed over the leader stage control circuitry that is located on the bottom of the amplifier board, was removed because it caused an asymmetry in the positive half of the current on the amplifier. The board pictured in Figure 5.23 originally had a shield and had been coated in super corona dope. The super corona dope did not flow into the area covered by the shield and an outline indicates where the shield use to be. In addition to the improved output current, not having the shield makes it significantly easier to replace Q10, Q12, and the optoisolators, which are one of the most likely components to need repair.

Chapter 6

Electrostatic trapping of OH

Collisions between cotrapped hydroxyl radical (OH) and rubidium present an interesting study of cold molecule and ultra-cold atom interactions. The cotrapped hydroxyl-rubidium (OH-Rb) system has the potential for sympathetic cooling of OH (hundreds of millikelvin) by rubidium (hundreds of microkelvin). Sympathetic cooling is the process of making one sample cold by placing it in thermal contact with a different sample that is even colder. It is a technique that has been used with ions [178] and atoms [179]. Sympathetic cooling of neutral molecules by ultra-cold atoms is a goal within the cold molecule community, and could help make the transition from millikelvin to microkelvin for direct cooling techniques.

Sympathetic cooling will be possible only if elastic collisions dominate over inelastic collisions. Elastic collisions, where momentum is exchanged, will cool down a sample, but inelastic collisions, where internal states are changed, will knock trapped molecules into untrapped states. In order for efficient sympathetic cooling to occur, there should be at least 10-100 elastic collisions for every inelastic collision [180].

Theoretical calculations suggest that the OH-Rb system has elastic cross sections that could be large enough compared to inelastic collision for sympathetic cooling to occur; however, background gas limited lifetimes and significant inelastic cross sections could prevent sympathetic cooling [181, 182]. These calculations have also indicated that optical or microwave dipole traps and alternating current traps would be more suitable for observing sympathetic cooling of OH by rubidium than an electrostatic trap, since the lowest energy state can be trapped. By trapping the

lowest energy state, inelastic collisions are limited since there is no lower energy state to fall into. They find that improved candidates for molecular sympathetic cooling should have high centrifugal barriers to suppress inelastic collision channels, a weak coupling of the electron to the internuclear axis (*ie.* Hund's case (b)), and a closed-shell atom to produce more isotropic potential energy surface [181, 182]. This suggests that a system with Li, Sr, or Ca and NH or CaH maybe better suited for sympathetic cooling. The CaF-Li and CaF-Rb systems have been modeled and shows promise toward sympathetic cooling down to 100 μK [183].

Nevertheless, the OH-Rb system is interesting to study due to the behavior of OH in electric and magnetic fields. OH has both an electric dipole moment (1.67 Debye) and a magnetic dipole moment ($2\mu_b$, where μ_b is the Bohr magneton). This gives rise to complex and interesting behavior in the presence of crossed electric and magnetic fields. This is opposed to a molecule such as ammonia, which has almost no magnetic dipole moment and thus would not have strong interactions with a magnetic field. OH exhibits both Stark and Zeeman effects and can be controlled with fields achievable in the laboratory. In the presence of combined fields, a complex energy landscape emerges with many avoided crossings that may shift significantly as the strength of the fields and the crossing-angle changes.

We proposed to study cotrapped OH and rubidium using a quadrupole electrostatic trap for OH and a quadrupole magnetic trap for rubidium. One goal will be to measure collision cross sections between electrostatically trapped OH and magnetically trapped rubidium. Although collision cross sections may not be favorable for sympathetic cooling, experimental measurements are still important, since small changes to the potential energy surfaces can lead to larger changes in the collision cross sections, making accurate predictions very challenging.

This study will compliment the many theoretical and experimental studies of cold OH that have already been performed. This includes theoretical studies of cold OH collision in a magnetic field [184] and a slow helium atomic beam colliding with trapped OH [185]. This second system has also been studied experimentally [92]. Cold OH has also been used in studies of the measurement of fundamental constants [10] and precision spectroscopy [186]. The effects of energy level crossings

cause by electric and magnetic fields on trapped OH have been observed experimentally in magnetic traps with electric bias fields and overlaid magnetic and electric quadrupoles [34, 187, 188].

In any study of trapped molecules, nonadiabatic losses are a concern. In the Stark deceleration calculations and experiments we have conducted thus far in this dissertation, we have made the assumption that the potential energy of the molecules is dependent only on the electric field strength and that the molecule remains in one quantum state. We take it for granted that as the fields change, the molecule reorients itself to follow the new axis and the potential that it is traveling on also changes smoothly. However, if the molecule changes quantum state at some point, the potential that it experiences would also change. Transitions to different quantum states are likely to occur if the rate of quantum state energy change is fast compared to the time it takes to traverse the region with small energy splitting. If a state change occurs, this can lead to transitions to a strong-field-seeking state and immediate trap loss or to more weakly trapped weak-field-seeking states for eventual trap loss. If the trap depth becomes shallower due to the new quantum state having a different Stark or Zeeman shift, molecules with large kinetic energies will be lost.

Studies of nonadiabatic loss of molecules from a trap have been completed using ^{12}CO and ^{13}CO [122] and for ammonia in an Ioffe-Pritchard type electrostatic trap [189]. Studies of nonadiabatic transitions inside a Stark decelerator found that losses are small for the decelerated group of molecules [190]. In this same study, they calculated that nonadiabatic transitions in decelerated OH are unlikely due to the large Λ -doubling splitting between the strong-field and weak-field seeking states and the linearity of the Stark shift at small fields.

In this chapter, we present the beginnings of an experiment working toward an OH-Rb co-trapped system, which will measure elastic and inelastic collision cross sections. We have measured electrostatically trapped OH distributions and discuss considerations regarding electrostatically trapped OH in a uniform magnetic field. Although the final OH-Rb trap configuration will have OH in a quadrupole electrostatic trap and Rb in a quadrupole magnetic trap, examining OH in a quadrupole electrostatic trap with a uniform magnetic field is a step toward building up the full experiment and model for the OH-Rb system. Additionally, it provides the opportunity to explore

configurations could potentially improve the density of trapped OH.

6.1 OH experimental setup

In order to slow down and trap OH molecules, we used our pin decelerator operating at ± 12 kV and a four electrode electrostatic trap at ± 10 kV. While the experiment operates in much the same fashion as described in Section 2.5, a few modifications have been made to accommodate the creation of OH and to slightly increase the deceleration and detection efficiency. OH is created via a discharge of 1% H₂O seeded in krypton. The gas mixture is pulsed into the chamber using a PZT valve with discharge plates mounted on the end, Figure 6.1. The discharge is formed between the two stainless steel plates (gray) which are separated by Macor disks (white). The resulting supersonic expansion produces a beam that is internally and translationally cold. The OH molecules are primarily in their rovibronic ground state with a mean longitudinal velocity of ~ 415 m/s. The pin decelerator is used to slow down the molecular packet to a mean velocity of 34 m/s. The detection region (Figure 6.2) has an electrostatic trap and two detection lasers for 1+1' REMPI of OH. The four trap electrodes are shown in blue and the two detection lasers, at 281 nm and 118 nm, counterpropagate between the two central electrodes.




Figure 6.1: The gas mixing system, valve, and discharge plates for creating a supersonic beam of OH. This figure is modified from Figure 3.16 in [155].




Figure 6.2: The detection region configuration for detecting trapped OH includes the last two sets of pins in the decelerator (gray), four trap electrodes (blue), 281 nm beam (red), and 118 nm beam (green). Taken from [153].

The pulsed-pin decelerator is operated in a hybrid mode that has presumably better transverse confinement and is referred to as *mixed deceleration*. In this method, a select number of $S=3$ deceleration stages are mixed into a timing sequence consisting of predominantly $S=1$ deceleration stages. The number of $S=1$ deceleration stages is chosen such that the deceleration phase angle is near $\phi_0=50^\circ$ and that the total number of remaining unused stages is even. $\phi_0=50^\circ$ is chosen here because due to transverse-longitudinal coupling, higher phase angles can lead to improved decelerator efficiency up to a phase angle near $\phi_0=50^\circ$ [140]. The unused stages are then spliced into the $S=1$ sequence to form $S=3$ deceleration stages. One stage in a $S=3$ deceleration stage is three times longer than an $S=1$ stage. Thus, there must be an even number of unused stages in order to extend a sequence from $S=1$ to partially $S=3$, since two additional stages are required per switch time. In the mixed mode, the $S=3$ deceleration stages are used as every third stage until they are all used up. The sequence is then followed by all $S=1$ deceleration stages until the end of the decelerator.

The sequence is typically calculated by selecting the number of desired $S=3$ stages and then checking the phase angle for the remaining number of deceleration stages. Figure 6.3 shows the time each stage lasts as a function of the stage number for one mixed deceleration timing sequence.

This sequence is for a timing molecule with an initial velocity of 447 m/s, a final velocity of 108 m/s, and a phase angle of $\phi_0 = 54.9^\circ$. Here, 20 S=3 stages are used. This results in a total of 109 deceleration stages of which 20 are S=3 and 89 are S=1. The underlying curve, without the spikes, is the S=1 sequence for 109 stages. Each stage lasts longer than the previous stage since the timing molecule slows down and takes longer to traverse subsequent stages. Every third stage for the first 60 stages is an S=3 deceleration stage and has been extended to be three times longer than the typical S=1 stage. This operation mode was recommended by David Reens and Hao Wu and yields a factor of two more signal. The success of this sequence was determined experimentally and careful 3D molecular dynamics simulations have not yet been performed to characterize the dynamics that cause this increase.




Figure 6.3: The length in time versus the stage number for one mixed deceleration timing sequence. This sequence has 20 S=3 stages and an initial velocity of 447 m/s, a final velocity of 108 m/s, and a phase angle of $\phi_0 = 54.9^\circ$.

After the OH molecules are decelerated, they are loaded into an electrostatic trap. The flexibility of the four electrode trap allows for various trap loading schemes. The Stark potential hills in the ‘half on’ and ‘trap on’ (Figure 2.20) configurations of the trap can be used as additional deceleration hills after the decelerator and before turning on the trap [155]. Figure 6.4 shows two examples of trap loading schemes. The optimized loading scheme (blue) was constructed so that the timing molecule is brought to a stop at the center of trap. The non-ideal loading scheme

(orange) does not bring the timing molecule to a stop, instead the timing molecule is located part way up the longitudinal potential energy hill when the trapping configuration is turned on for the last time. This results in greater average energy of molecules in the trap, which is normally not desired. However, there may be certain experiments where trapped molecules with greater energy are desired (see Section 6.4).

With ± 10 kV on the trap electrodes, the trap depth for OH is 700 mK in the z direction and 250 mK in the x and y directions. The molecules are held in the trap for a set amount of time before they are released and detected via a 1+1' REMPI scheme. To detect the OH molecules, the 281 nm photon is created by a dye laser and the 118 nm photon is created by frequency tripling 355 nm light from a Nd:YAG (Section 6.5). The 281 nm beam is focused using a cylindrical lens so that it can fit between the two central electrodes of the trap and is uniform across the vertical (y) axis of the trap. The 118 nm beam is focused on the detection region using a $f=25$ cm lens. Two picomotors are mounted to the last steering mirror for the 355 nm beam. This allows for both horizontal and vertical movement of the 355 nm beam, which in turn steers the 118 nm beam. By moving the 118 nm beam, the vertical density profile of the entire trapped sample can be mapped out. The 281 nm beam is approximately uniform over the entire trapped sample so that there is a fairly uniform extraction efficiency vertically across the trap. On the other hand, the horizontal density profile cannot be easily detected since the 118 nm beam has limited mobility in the horizontal direction due to the trap electrodes. Additionally, the extraction efficiency of the TOFMS is dependent on the horizontal position of the detection laser making the extraction non-uniform in this direction.

Once the OH molecules have been ionized, they are accelerated to the MCP detector using a TOFMS configuration of [2,2,1,0] kV on the four trap electrodes. This is slightly different than previous work in our lab [95,155,191], which always used [1,1,0,0] kV. It was found that the [2,2,1,0] kV configuration give a more uniform extraction efficiency as the 118 nm beam position is scanned in the vertical direction.

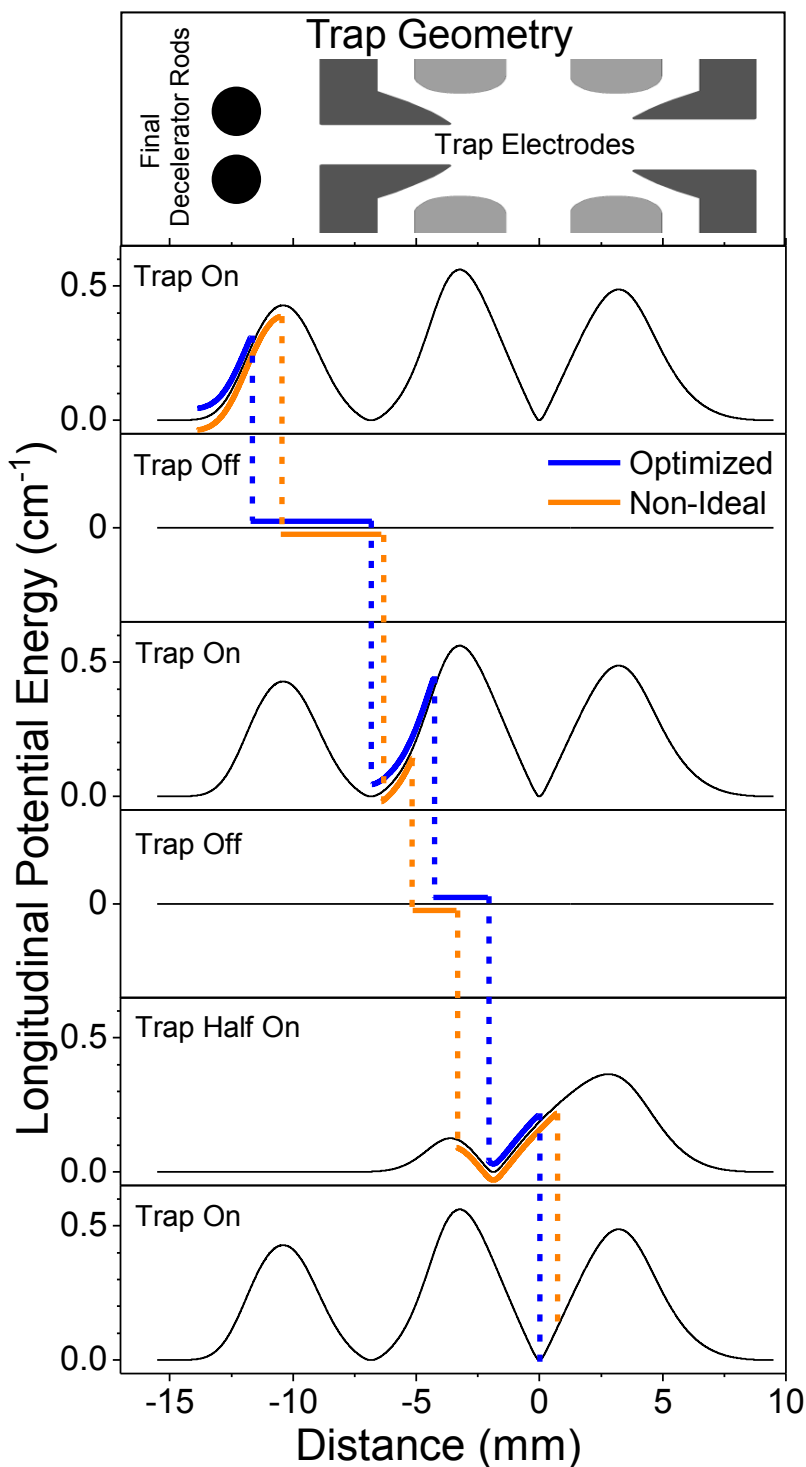


Figure 6.4: A diagram of two different OH trap loading schemes. The top panel shows the location of the last pair of pins in the decelerator and the four trap electrodes. The subsequent panels illustrate two trap loading schemes. Blue (optimized) and orange (non-ideal) lines show the location of the timing molecule when the potential is switched from one configuration to another. Both schemes use the same sequence of configurations: on, off, on, off, half on, and on, but switch at different times. Figure taken from [153]

6.2 OH REMPI Detection

OH has only four bound electronic states: $X^2\Pi$, $A^2\Sigma^+$, $B^2\Pi$, and $C^2\Sigma^+$ [192]. The $B\leftarrow X$ and $C\leftarrow X$ transitions lie in the vacuum ultraviolet (VUV) range where there are few tunable sources of laser light. Thus the $A\leftarrow X$ transition is the most promising one for state selective detection of OH. OH has an ionization potential of 13.017 eV. While it has been detected using 2+1 [193, 194] and 3+1 REMPI schemes [195], the observed ion signals have been weak. LIF on the $A^2\Sigma^+ - X^2\Pi$ transition is readily used in many OH experiments [87, 101, 108, 188], but as discussed in Section 2.5.2, REMPI provides the added ability to map out the density distribution of trapped molecules.

The $A^2\Sigma^+$ excited-state electron configuration of OH is $(1\sigma)^2(2\sigma)^2(3\sigma)^1(1\pi)^4$, while the ground-state electron configuration of OH^+ is $(1\sigma)^2(2\sigma)^2(3\sigma)^2(1\pi)^2$. This suggests that a one-photon process between these two states should be forbidden since two electrons must change orbitals. However, a 1+1' REMPI process, which makes use of an autoionizing Rydberg state that converges on a higher ionization limit, has been experimentally demonstrated by Beames *et al.* [196]. This scheme has subsequently been used in other OH studies [152, 153, 197]. This REMPI scheme makes use of one 281 nm and one 118 nm photon. The 281 nm photon is typically generated via a doubled dye laser. The tunability of a dye laser is useful in this case so that different excited states can be used as the intermediate. The second, 118 nm, photon is generated by frequency tripling the third harmonic (355 nm) of an Nd:YAG in a gas cell. This procedure is discussed in detail in Section 6.5.

Figure 6.5 shows the relevant OH energy levels for REMPI detection with the transitions used in this experiment in bold. From a strictly energy stand point, 118 nm light is able to ionize OH from any rotational level in the excited $A^2\Sigma^+, v = 1$ manifold. However, careful choice of the excited state will make a significant difference in the resulting ion signal. The trapped OH molecules are in the upper Λ -doublet of the rotational and vibrational level ($J=3/2, v=0$) of the ground $X^2\Pi_{3/2}$ electronic state. The strongest $A \leftarrow X$ LIF transition is Q_1 (1.5), but ionization from the ($J=3/2, p=-$) state yields minimal ion signal. In a supersonic beam of OH, the largest




Figure 6.5: An energy level diagram of OH showing the relevant states for 1+1' detection. The energy levels are not to scale. Q_1 (1.5) is the typical LIF transition, Q_{12} (1.5) is the strongest transition between the $X^2\Pi_{3/2}$ and $A^2\Sigma^+$ states, and R_{12} (1.5) is the transition used in the 1+1' REMPI scheme for OH. Figure taken from [153].

ionization signal is achieved on the $Q_{12}(1.5)$ transition [152]. However $Q_{12}(1.5)$ originates from the lower Λ -doublet, which is not decelerated and trapped. Instead ionization from the R_{21} (1.5) transition is used in this work. A 118 nm photon from the ($J=5/2$, $p=-$) state is resonant with the $A^3\Pi$, 3d autoionizing state which helps to enhance the ion signal.

6.3 OH 1+1' REMPI with 118 nm photons

The advantage of using REMPI detection over LIF is the ability to scan the laser over the sample and map out the density distribution using a setup such as the one described in Section 6.1. Although decelerated OH is commonly detected using LIF [75, 108, 116, 188], the 1+1' REMPI scheme opens the doors to improved signal-to-noise and data analysis methods. There are many limits to 118 nm light generation, but enough 118 nm power can be created to make the difficulties of setting up a REMPI scheme for OH worth it. Below, we show a few measurements that can only be performed by scanning the detection beam and would not be possible using only LIF since LIF

requires detecting the entire trapped sample at once. An expanded beam, which illuminates the entire trapped sample, is used in LIF. While the beam could be made smaller, the signal-to-noise become impractically low and a localized LIF beam is not a practical detection method.

In the case where two different trap loading schemes have been used (see Figure 6.4), and each results in different density distributions in the trap [153], differences in the central density and width of the trapped molecules can be observed via 1+1' REMPI. These measurements would not be possible with LIF detection where the entire sample is illuminated and detected at once. Figure 6.6 shows measurements of the central OH density in both the experiment and simulations. In the optimized loading scheme, where the timing molecule is stopped at the center of the trap, more molecules are found near the center of the trap. On the other hand, in the non-ideal case, where the trap turns on when the timing molecule is part way up one side of the longitudinal potential, there is lower density at the center and many more molecules orbit at large radii about the center of the trap. The density distribution at different times can be used to determine the full-width half-maximum (FWHM) of a trapped sample over time, Figure 6.7. Larger changes in the width can be seen in non-ideal loading (Figure 6.7b) than in optimized loading (Figure 6.7a). This is expected since, on average, the molecules in the non-ideal loading scheme have a greater residual oscillations.

The FWHM of the cloud the vertical (y) direction is determined by recording vertical density scans at set detect times and then fitting the distributions. The FWHM is then extracted by fitting these scans. 3D Monte Carlo molecular dynamics simulations of molecules in the trap produce similar width modulations as seen in the experiment. Such modulations can be used to diagnose issues with the detection laser alignment or with the trap loading scheme. Since trap loading sequences are determined based on the final velocity of the timing molecule, if the resulting decelerated molecule packet does not have the same average velocity and position as the timing molecule, the resulting molecule signal in the trap can be very low. By iterating through trap loading sequences or adjusting the position of the 118 nm beam, it can be determined whether the trap was loaded well or not.




Figure 6.6: Experimental (solid) and simulated (dashed) measurements of the central density OH in the trap for the two different loading schemes depicted in Figure 6.4. (a) the optimized loading scheme shows smaller oscillations when loading and an overall greater density at the center of the trap once the oscillations damp out. (b) the nonideal loading scheme has larger loading oscillations and a lower central density once the oscillations damp out. The difference in the central density can be seen in phase-space plots of the simulations. In the optimized case, the majority of molecules are at the center of the trap. In the nonideal case the molecules are primarily at large orbits about the center of the trap. Figure from [153].

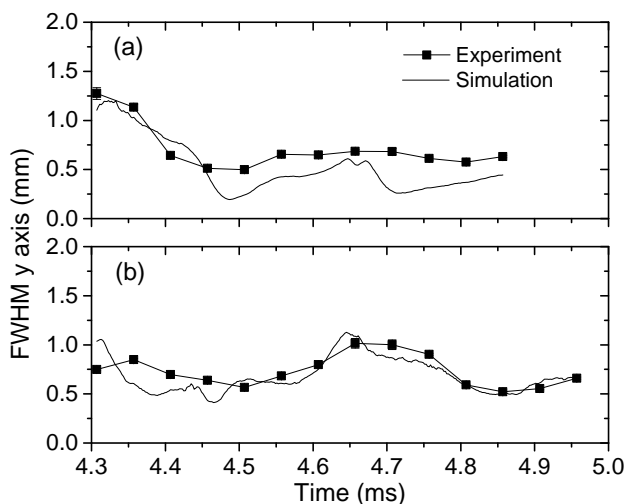


Figure 6.7: The FWHM of trapped OH over time for a) optimized loading and b) non-optimized loading. Both simulations (solid lines) and experimental measurements (squares) are shown. In both cases, the simulations reproduce the modulation of the cloud after it is trapped as seen in the experiment. The statistical uncertainty is on the order of the point size or smaller. Figure from [153].

Lastly, the ability to map out the density distribution allows for the measurement of the

effective temperature of the trapped molecules. The density profile can be measured at various times after the trap is turned off. From these measurements, the FWHM can be extracted and the change in width used to determine a temperature. Figure 6.8 shows one of these measurements. Here, OH was trapped using the optimized loading scheme and held for 15 ms before being released from the trap. A fit to this expansion gives a V_{rms} (root-mean-square velocity) of 5.5 ± 0.2 m/s and 6.2 ± 0.2 m/s for the experimental and simulated results respectively [153].

The difference in the initial size of the cloud in the simulation and experiment can likely be attributed to uncertainty in the initial distribution of molecules in the trap. In the simulations, an assumption is made about the initial distribution in a Stark deceleration experiment, particularly that its position and velocity distribution are Gaussian. The Gaussian molecular packet propagates through the decelerator and trap loading potentials. Deviations between the packet in the experiment compared to the simulation will propagate through the simulations and lead to differences between the observed molecule signal and the simulated signal. It is important to remember here that this is not a true temperature since the trapped molecules do not have enough collisions to thermalize. Instead, this is a measurement of the average energy of the trapped molecules.

These measurements in Figures 6.6, 6.7, and 6.8 demonstrate some of the ways in which the dynamics of trapped molecules can be observed and characterized. While these dynamics have often been investigated using simulations, they are not often observed experimentally. This method of observing OH molecules using 1+1' REMPI demonstrates the advantage of using an ionization detection technique even though it is more complicated than LIF, since it requires two lasers.




Figure 6.8: The width of the OH cloud as it expands after the trap is turned off. Figure from [153].

6.4 OH in electric and magnetic fields

With a reliably operating electrostatic trap and a REMPI detection scheme that allows us to map the vertical density distribution in the trap and fully understand and characterize the trapped OH molecules, we can now examine trapped OH under the influence of other fields and particles. To perform cotrapped OH-Rb experiments, the ultra cold rubidium will be prepared separately in a rubidium cell that is attached to the trap chamber. A MOT of rubidium is loaded into a magnetic trap. The coils for the magnetic trap are mounted on a translating track. The magnetically trapped rubidium is then transported from the rubidium cell into the trap chamber of the decelerator experiment and centered on the electrostatic trap. This procedure is an almost identical to the one used in [94], but with a slightly different geometry.

Given this coil setup, we also have the ability to apply a uniform magnetic field to the electrostatically trapped OH. The two coils, which are external to the vacuum chamber, can be placed in a Helmholtz configuration to apply a uniform magnetic field to the trap. Figure 6.10 shows this setup where a uniform magnetic field is applied along the y axis. The coils have 24 turns each and can carry a current up to 500 A. In the Helmholtz configuration, this gives a maximum field of 800 G at the center of the trap. Figure 6.9 shows a plot of the electric fields in the longitudinal

and transverse directions with ± 10 kV on the trap electrodes. Near the bottom of the trap, the gradient in the longitudinal and transverse directions are 123 kV/cm^2 and 62 kV/cm^2 respectively. This configuration will be used to study electrostatically trapped OH in a uniform magnetic field.




Figure 6.9: The electric field in the trap along (a) a line along the z axis and (b) a line along the y axis. The asymmetry in the longitudinal direction is due to the grounded decelerator electrodes that are upstream of the trap.




Figure 6.10: The electrostatic trap with Helmholtz coils to produce a uniform magnetic field at the center of the trap. While the trap and decelerator are located in the UHV chamber, the coils are external to the chamber. This is almost the same setup as in Figure 6.2, but with the addition of the coils.

6.4.1 Hamiltonian for OH in electric and magnetic fields

With the ability to apply a uniform magnetic field, we must now examine how the OH spin-orbit and rotational states behave in crossed electric and magnetic fields. The strong coupling of \mathbf{L} and \mathbf{S} in Hund's case (a) means the internuclear axis will want to follow both the magnetic and electric field lines. Here we summarize the treatment of the ${}^2\Pi_{3/2}$ ground state of OH in crossed fields.

The same convention as in [198] is used, the magnetic field B is restricted to one axis and the electric field E forms an angle β with B . The Hund's case (a) basis $|J, M, \bar{\Omega}, \epsilon\rangle$ is used. The Hamiltonian for the ground state of OH in crossed electric and magnetic fields is [34, 199]

$$\begin{pmatrix} -\frac{\Delta}{2} - \frac{6}{5}\mu_b B & 0 & 0 & 0 & \frac{3}{5}\mu E \cos \beta & -\frac{\sqrt{3}}{5}\mu E \sin \beta & 0 & 0 \\ 0 & -\frac{\Delta}{2} - \frac{2}{5}\mu_b B & 0 & 0 & -\frac{\sqrt{3}}{5}\mu E \sin \beta & \frac{1}{5}\mu E \cos \beta & -\frac{2}{5}\mu E \sin \beta & 0 \\ 0 & 0 & -\frac{\Delta}{2} + \frac{2}{5}\mu_b B & 0 & 0 & -\frac{2}{5}\mu E \sin \beta & -\frac{1}{5}\mu E \cos \beta & -\frac{\sqrt{3}}{5}\mu E \sin \beta \\ 0 & 0 & 0 & -\frac{\Delta}{2} + \frac{6}{5}\mu_b B & 0 & 0 & -\frac{\sqrt{3}}{5}\mu E \sin \beta & -\frac{3}{5}\mu E \cos \beta \\ -\frac{3}{5}\mu E \cos \beta & -\frac{\sqrt{3}}{5}\mu E \sin \beta & 0 & 0 & \frac{\Delta}{2} - \frac{6}{5}\mu_b B & 0 & 0 & 0 \\ -\frac{\sqrt{3}}{5}\mu E \sin \beta & \frac{1}{5}\mu E \cos \beta & \frac{2}{5}\mu E \sin \beta & 0 & 0 & \frac{\Delta}{2} - \frac{2}{5}\mu_b B & 0 & 0 \\ 0 & -\frac{2}{5}\mu E \sin \beta & -\frac{1}{5}\mu E \cos \beta & -\frac{\sqrt{3}}{5}\mu E \sin \beta & 0 & 0 & \frac{\Delta}{2} + \frac{2}{5}\mu_b B & 0 \\ 0 & 0 & -\frac{\sqrt{3}}{5}\mu E \sin \beta & -\frac{\sqrt{3}}{5}\mu E \cos \beta & 0 & 0 & 0 & \frac{\Delta}{2} + \frac{6}{5}\mu_b B \end{pmatrix}. \quad (6.1)$$

The state ordering from left to right and top to bottom begins with $|J, M, \bar{\Omega}, \epsilon\rangle = |\frac{3}{2}, -\frac{3}{2}, \frac{3}{2}, e\rangle$ and ends with $|\frac{3}{2}, \frac{3}{2}, \frac{3}{2}, f\rangle$, where first M is increased and then ϵ .

The fully stretched states do not experience any actual or avoided crossings in the crossed fields. At large \mathbf{E} or \mathbf{B} fields the states approach the case for \mathbf{E} -field only or \mathbf{B} -field only. However, at intermediate values of \mathbf{E} and \mathbf{B} with certain values of β , there are level crossings and avoided crossings. Additionally, in combinations of strong electric and magnetic fields neither field provides a suitable quantization axis. However good quantum numbers can still be determined using a 'Hund's case-X' [171].

The Hamiltonian in Equation 6.1 can be used for the uniform magnetic field case discussed above since the magnetic field, \mathbf{B} is always in the same direction. The changing magnitude of the electric field is represented by \mathbf{E} and the changing direction of the electric quadrupole with respect to the y-axis is represented by β .

Figures 6.11-6.13 each show how the various states of the $^2\Pi_{3/2}$ ground state of OH evolve under different magnetic and electric fields. Each figure has a constant magnetic field of 90, 500, and 800 Gauss. These values correspond with 50, 300, and 500 A in our coils. 50 and 500 A are achievable in our setup, however, 300 A is not due to power dissipation limitations on the current driving MOSFETs. Despite this, 500 G (300 A) has been plotted to give a better sense of how the energy surfaces evolve by changing angle and field. The electric field range on the x axis covers the range found in the center of quadrupole electrostatic trap. The maximum electric fields in the trap are 30 kV/cm and 16 kV/cm in the longitudinal and transverse directions respectively. Within each figure, β , the angle between \mathbf{E} and \mathbf{B} , takes on the values of 0, $\pi/16$, $\pi/8$, $\pi/4$, $3\pi/8$, and $\pi/2$ starting from the top left and moving across each row.

Small uniform magnetic fields (*ie.* B=90 G, Figure 6.11) do not significantly change the trap depth at any angle. However, crossing and avoided crossing between the various states Zeeman and Stark shift states begin to emerge. At large uniform magnetic fields (*ie.* B=800 G, Figure 6.13), the maximally stretched state $|\frac{3}{2}, \frac{3}{2}, +\frac{3}{2}, f\rangle$ (blue) changes slope with changing angle, β . This means that trap depth will be modified with a large uniform magnetic field.

6.4.2 Electrostatically trapped OH with a uniform magnetic field

As the molecules traverse the trap in the presence of a strong magnetic field, they will encounter a different potential energy landscape than in the no-field case. If they are in a state that has an avoided or actual crossing with another state, they will follow either adiabats or diabats depending on the energy gap and molecule velocity. This will alter either the trap depth that the molecule encounters or it may cause the molecule to fall into an untrappable state.

In the case of the electrostatic trap with a strong uniform magnetic field, the trap depth is slightly weaker than in the purely electrostatic trap. Molecules out of the decelerator will be in the $|f; +\frac{3}{2}\rangle$ state. Once in the trap, they should predominantly remain in the $|\frac{3}{2}, \frac{3}{2}, +\frac{3}{2}, f\rangle$ state. As a quick estimate, we take the case where the minimum energy gap between the $|\frac{3}{2}, \frac{3}{2}, +\frac{3}{2}, f\rangle$ state and any other state is less than 10 MHz to have any probability of a diabatic transition. This only

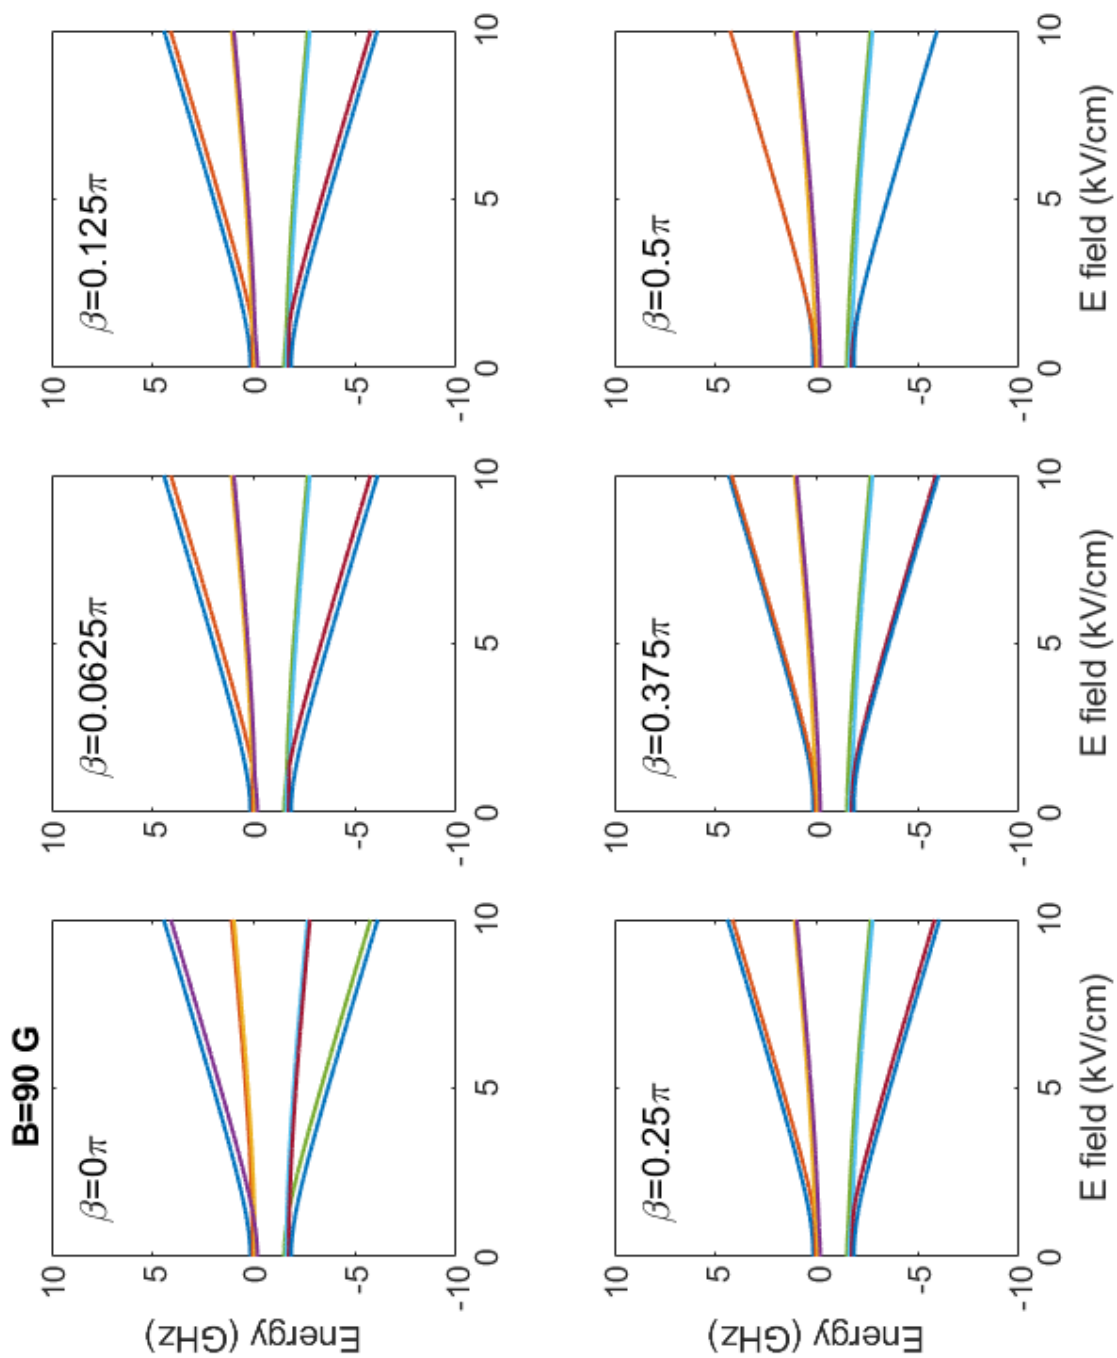


Figure 6.11: OH energy levels in a uniform 90 G magnetic field with varying β and electric field.




Figure 6.12: OH energy levels in a uniform 500 G magnetic field with varying β and electric field.




Figure 6.13: OH energy levels in a uniform 800 G magnetic field with varying β and electric field.

occurs at electric fields greater than 6 kV/cm and $\beta = \pi/2$ (See Figure 6.13). In our trap, these two conditions occur only in the x-z plane and far away from the trap center. Figure 6.14 shows a cut in the x-z plane of our trap with the regions where the electric field is greater than 6 kV/cm colored in red. This create a “zone of death” that is a sheet, but with a hole at the center of the trap.




Figure 6.14: The region in the x-z plane ($y=0$) of our trap where the electric field is greater than 6 kV/cm is shown in red. For a uniform field of 800 G in our electrostatic trap this is the only plane in which trapped OH molecules encounter any avoided crossings with another fine structure state and have a chance of falling out of the trap.

This is compared to the case of no magnetic field, where there is a degeneracy between the $M_J\Omega = -9/4$ and $-3/4$ states when $\mathbf{E}=0$ at the center of the trap. Molecules traveling on trajectories that come close to the trap center have a chance of crossing from the $M_J\Omega = -9/4$ state to the $M_J\Omega = -3/4$ state. In the $M_J\Omega = -3/4$ state, molecules are more weakly trapped. This can lead to trap loss if the kinetic energy of the molecule can overcome the trap potentials. Over the course of the background gas limited trap lifetime (1-2 s), the molecule may traverse this crossing hundreds of times. Taking the same approach as in [188], the transition probability when the molecule crosses through an energy minimum between two states can be calculated from the Landau-Zener hopping probability, P ,

$$P = e^{-\pi a^2 / 2\hbar \dot{G}}, \quad (6.2)$$

where a is the coupling between the two states in question and \dot{G} is the rate of approach. In one

dimension only, \dot{G} can be found via

$$\dot{G} = v_z \frac{dG}{dz}. \quad (6.3)$$

Simulations in which multiple molecules evolve on more than one potential energy surface and the crossings between $M_J\Omega = -9/4$ state to the $M_J\Omega = -3/4$ state have been calculated using Equation 6.2 are needed in order to understand these trap losses in both the electrostatic trap with no magnetic field. Once completed these simulations will be compared with measured experimental trap life times for both the uniform magnetic field case and no field case.

6.4.3 Future work

In the uniform magnetic field case, it appears that trap losses will occur only at large radii. If a molecule distribution in the trap has a greater average energy, with many molecules orbiting at large radii, it may experience greater trap loss than one that has a lower average energy, with the majority of molecules orbiting near the trap center. We have already demonstrated that two trapped samples such as this can be prepared and measured (Figure 6.6) [153]. Poor trap loading will lead to larger average energies in the trap and molecules that orbit the trap at larger radii. If molecules can be intentionally loaded into the ‘zone of death,’ trapped molecules in the non-ideal loading scheme will have shorter trap lifetimes.

After characterizing the electrostatic trap with a uniform magnetic field, the electrostatic trap with a quadrupole magnetic trap and OH-Rb collisions will be examined. For the OH-Rb system, there is a quadrupole electric trap with a quadrupole magnetic trap, and in this case there will not be perpendicular electric and magnetic fields. However, since the gradients between the two traps are different, the angle between the electric and magnetic fields (which will be $< 90^\circ$) will still change. A similar electric and magnetic field model to above will still be needed to study this system. The Hamiltonian for the system cannot be taken to just be in Equation 6.1, since the magnetic field is not uniform along one axis. Instead, a lab frame Hamiltonian will be required. Once this has been characterized elastic and inelastic collision cross sections can be measured in the OH-Rb system.

6.5 118 nm light

118 nm (10.5 eV) photons fall in the vacuum ultraviolet (VUV) range. While VUV photons are very useful as ionization sources in mass spectrometry, few table-top high flux VUV sources exist. The majority of VUV light used in experimental physical chemistry comes from synchrotrons, but the complexities, scale, and location of these sources limits the number of experiments that can use these sources. When generating table top VUV, either an excimer laser or a gas cell containing a nonlinear media is used. VUV light cannot be generated inside a crystal because crystals become opaque around 200 nm.

To create the 118 nm light required for OH REMPI detection, we triple the third harmonic output of a Nd:YAG (355 nm) in a xenon-argon gas cell. 118 nm is the third harmonic of 355 nm, which is an easy to obtain wavelength from commercial YAG lasers. A tripling process is chosen since the isotropic nature of gases limits the nonlinear process to only odd-orders. In order for even numbered processes to occur in a gas, a symmetry breaking field such as a strong electric or magnetic field must be applied. At 118 nm, xenon is negatively dispersive and xenon acts as the tripling medium. Xenon is the only rare gas that has a negative dispersion at 118 nm. Argon is positively dispersive at 118 nm and acts as the phase matching medium. The phase-matching gas is required so that the fundamental light (355 nm) and third-harmonic light (118 nm) maintain a fixed phase relationship and 118 nm light created at slightly different locations build coherently. While krypton or neon could also be used as the phase-matching gas, as will be seen below, krypton is less efficient than argon, and neon requires significantly higher total pressures, which adds experimental challenges.

The 118 nm photon can be used in a single or multi-photon ionization molecular detection technique. As seen in Section 6.2, 118 is well suited for 1+1' REMPI of OH due to the auto ionizing state that can be accessed from the $A^2\Sigma^+, v = 1$ excited state. Additionally, for molecules with a lower ionization potential, 118 nm photons can be used for single photon ionization. This is especially useful when constructing a 118 nm detection system, since it can be tested using a

molecule such as ammonia or propyne, which have ionization potentials below 10.5 eV. The xenon-argon 118 nm system has been used for various molecule detection schemes including hexane [200], methanol [201], and OH [152, 196].

The sensitivity of the 1+1' REMPI detection technique for OH is limited by the flux of 118 nm photons, since there is significantly lower 118 nm power than the first 281 nm photon. The conversion efficiency from 355 nm to 118 nm is low and an additional suppression factor occurs at moderate xenon pressures (>20 torr), which limits the overall efficiency of this detection method. Here, we give an overview of 118 nm light production and explore the suppression at higher pressures. More details on the mechanisms and calculations for the tripling process can be found in [155, 200, 202, 203].

6.5.1 Creating 118 nm light

The power at the third-harmonic, $\mathcal{P}_{3\omega}$, for an incident fundamental power \mathcal{P}_ω , where both powers are in watts, is given by [202, 203]

$$\mathcal{P}_{3\omega} = \frac{8.215 \times 10^{-2}}{(3\lambda)^4} \left| \chi_{eff}^{(3)} \right|^2 \mathcal{P}_\omega^3 |F_1(b\Delta k)|^2. \quad (6.4)$$

Here λ is the third-harmonic wavelength in cm and $\chi_{eff}^{(3)}$ is the macroscopic nonlinear coefficient. $F_1(b\Delta k)$ is a dimensionless function, Δk is the wave-vector mismatch between the fundamental and third-harmonic, and b is the confocal parameter. The confocal parameter is a measure of how tightly a beam is focused. It is equal to twice the Rayleigh length. Since we are discussing only 118 nm light production, λ is set and there are three free parameters left that can be varied experimentally: the peak powers in the fundamental, the macroscopic nonlinear coefficient $\chi_{eff}^{(3)}$, and the phase matching function $F_1(b\Delta k)$.

The peak power in the fundamental is limited only by the output power of the Nd:YAG laser and the resulting power density on any optics in the beam path. Thus, to create as much 118 nm light as possible, \mathcal{P}_ω should be as large as possible. Independent maximization of the last two parameters is impossible because both $\chi_{eff}^{(3)}$ and $F_1(b\Delta k)$ are dependent on \mathcal{N} , the atomic number

density of the gas. $\chi_{eff}^{(3)}$ is given by

$$\chi_{eff}^{(3)} = \mathcal{N} \times \chi_a^{(3)}, \quad (6.5)$$

where \mathcal{N} has units of number/cm³ and $\chi_{eff}^{(3)}$ is in *esu*.

The last factor, $F_1(b\Delta k)$, is the phase-matching function. It is an integral over the reduced longitudinal coordinates of the gas cell and generally evaluated numerically. However, since the nonlinear medium is much longer than the confocal parameter of the focused fundamental beam ($b \ll L$), $F_1(b\Delta k)$ in this tight-focusing limit can be solved to give [203, 204]

$$|F_1(b\Delta k)|^2 = \begin{cases} \pi^2 (b\Delta k)^2 e^{b\Delta k} & \Delta k < 0 \\ 0 & \Delta k \geq 0. \end{cases} \quad (6.6)$$

Figure 6.15 show a plot of Equation 6.6 for $b = 1$. This function has has a maximum at $\Delta k(\lambda_0)_{max} = -\frac{2}{b}$. This means in order to maximize 118 nm light production, Δk should be adjusted based on the confocal parameter, which is set by the lens used to focus the 355 nm light into the gas cell.




Figure 6.15: Plot of $|F_1(b\Delta k)|^2$ in the tight-focusing limit and for $b = 1$. The function has a maximum at $\Delta k(\lambda_0)_{max} = -\frac{2}{b}$.

Δk in Equation 6.6 is related to \mathcal{N} through the index of refraction,

$$\Delta k = \frac{2\pi}{\lambda} [n_1 - n_3]. \quad (6.7)$$

In this notation, n_1 and n_3 are the index of refraction at the third-harmonic wavelength λ and the fundamental wavelength 3λ , respectively. For a gas, the effective macroscopic index of refraction,

n , is given by

$$n(\lambda) = 1 + \mathcal{N}\chi_a^{(1)}(\lambda), \quad (6.8)$$

where $\chi_a^{(1)}$ is the atomic linear susceptibility at wavelength λ . Now Δk can be related to \mathcal{N} and expressed as

$$\Delta k = C\mathcal{N}, \quad (6.9)$$

where

$$C = \frac{2\pi}{\lambda} [\chi_a^{(1)}(\lambda) - \chi_a^{(1)}(3\lambda)] \quad (6.10)$$

is a constant describing the microscopic dispersive properties of the gaseous medium. Δk , and thus $F_1(b\Delta k)$, can be optimized by simply varying the pressure of the gas.

As can be seen from Equation 6.6, Δk must be negative to generate third-harmonic light. This occurs only in media that exhibits anomalous dispersion (i.e., $n_3 > n_1$). Xenon was chosen as the nonlinear medium for THG of 118 nm because it is the only rare gas that exhibits anomalous dispersion at this wavelength. With only a negatively dispersive gas, Δk and $\chi_{eff}^{(3)}$ are both dependent on \mathcal{N} and $\chi_{eff}^{(3)}$ and $|F_1(b\Delta k)|^2$ in Equation 6.6 cannot be varied independently. If instead a positively dispersive gas with negligible nonlinearity is added, Δk and $\chi_{eff}^{(3)}$ can be decoupled. Argon is typically used as the positively dispersive gas for THG of 118 nm light. For a mixture of the two gases, the total wave-vector mismatch at the nominal third-harmonic wavelength λ_0 is given by summing over the partial pressures of each gas,

$$\begin{aligned} \Delta k(\lambda_0) &= \mathcal{N}_{Xe}C_{Xe}(\lambda_0) + \mathcal{N}_{Ar}C_{Ar}(\lambda_0) \\ &= \mathcal{N}_{Xe} [C_{Xe}(\lambda_0) + RC_{Ar}(\lambda_0)] \end{aligned} \quad (6.11)$$

where $R = \frac{\mathcal{N}_{Ar}}{\mathcal{N}_{Xe}}$ is the pressure ratio of the two gases. With the addition of argon, $\chi_{eff}^{(3)}$ (Equation 6.5) can be adjusted using \mathcal{N}_{Xe} , while Δk (Equation 6.11) can be adjusted using \mathcal{N}_{Ar} .

In order to maximize $F_1(b\Delta k)$ and find the optimum phase-matching ratio of the two gases in the tight-focusing limit, Eq. 6.11 must equal $-\frac{2}{b}$, the value of $\Delta k(\lambda_0)_{max}$ at the maximum of Eq. 6.6. After changing the number density, \mathcal{N} , to pressure, \mathcal{P} , via the ideal gas law, this gives an

optimum ratio of

$$\left(\frac{\mathcal{N}_{Ar}}{\mathcal{N}_{Xe}}\right)_{opt} = -\frac{2}{b} \frac{1}{\mathcal{N}_{Xe} C_{Ar}(\lambda_0)} - \frac{C_{Xe}(\lambda_0)}{C_{Ar}(\lambda_0)}. \quad (6.12)$$

For large xenon pressures, which are desirable for increasing $\chi_{eff}^{(3)}$ and thus $\mathcal{P}_{3\omega}$, the optimum ratio approaches the constant value of $\frac{C_{Xe}}{C_{Ar}}$ and is independent of the value of confocal parameter b . Using the values $C_{Xe}(118.2 \text{ nm}) = -6.12 \times 10^{-17} \text{ cm}^2$ and $C_{Ar}(118.2 \text{ nm}) = +5.33 \times 10^{-18} \text{ cm}^2$ from Mahon *et al.*, [205] this ratio is 11.48.

6.5.2 Experimental setup

The experimental setup for the gas cell used to create 118 nm light is shown in Figure 6.16. The gas mixing cell (tan) is ~ 40 cm long and is constructed of three 2-3/4" standard stainless steel ConFlat crosses. It is important that the cell diameter not be too small at any point. This restricts gas flow and leads to inconsistent gas mixing and variable 118 nm signal. For example, we found that when we attached a section of 1/4" tubing to the gas cell, xenon, the first gas placed in the cell would become trapped in the narrower section and diffuse out over time. This would significantly change the gas ratio as a function of time and decrease 118 nm light production. The input of the cell is sealed with a UV fused silica window. It is important to make sure that the power density of the 355 nm beam at the window is low enough that it does not burn. 355 nm light, produced by a Continuum Inlight III-10, which is internally tripled, is focused into the gas mixing cell with a 50 cm focal length lens. The second lens in the system is a 25 cm focal length MgF₂ lens and focuses the 118 nm beam in between the TOFMS electrodes. This lens sits in a custom conflat holder [155] so that it can also act as a seal between the gas mixing cell and the detection chamber. MgF₂ was chosen for this lens because of its VUV transmission and hardness, since many optics are not transmissive in the VUV. A Macor iris (not shown) is placed after the MgF₂ lens on the detection chamber side, which blocks the excess 355 nm light that was not converted to 118 nm light.

Once the 118 nm photons are created, they are difficult to detect since they cannot propagate through air. Here, we use ionization of a uniform molecular cloud in order to detect the 118 nm




Figure 6.16: A diagram of the 118 nm gas mixing cell and the ionization and detection region. The 335 nm beam (red) is focused into a gas cell containing xenon and argon (tan). 118 nm light is formed at the focus of the 355 nm light via THG if the gas cell contains the correct ratio of xenon to argon. The 118 nm light (purple) is focused into the detection chamber (light blue). Inside the detection chamber four electrodes (dark blue) accelerate ions formed at the focus of the 118 nm beam toward the MCP detector. The detection chamber can either be flooded with a low pressure gas or a decelerated beam can be loaded into the trapping region. The blue arrow in the detection chamber indicates the molecular beam axis of the decelerator. Slow OH molecules enter the trap along the arrow.

photons. This is easier than using a molecular beam since the arrival time of the molecules in the detection region is not a factor. This is accomplished by flooding the detection chamber with ND_3 , a molecule that can be ionized with a single 118 nm photon. Typically, the detection chamber is filled to a partial pressure of $\sim 9 \times 10^{-8}$ torr. Since the ND_3 pressure in the detection chamber is kept constant, the ND_3 ion signal acts as a proxy for the 118 nm photon flux. The number of ND_3^+ ions created will then be proportional to the number of 118 nm photons that reach the detection region. The resulting ND_3^+ ions are accelerated by a TOFMS to an MCP detector as described in Section 6.1.

6.5.3 Limitations to 118 nm light production

As seen in Equation 6.6, increasing the xenon pressure and the corresponding phase-matching argon pressure should result in more 118 nm light production. However, this is not the case experimentally. Figure 6.17 shows the 118 nm flux for various pressures of xenon gas and a constant

355 nm input power. Each colored curve shows a constant pressure of xenon gas with varying pressures of argon gas. The peak of each curve is the phase-matched mixture for the given pressure of xenon. For $P_{Xe} > 15$ torr the phase-matched pressure is 11.5 as predicted by Equation 6.12. Each curve is recorded by adding a set amount of xenon gas to the mixing cell, then argon is added slowly while monitoring the ionized ND₃ signal. As argon is added, the gas mixture will take time to mix. This may take 1-15 minutes depending on how the gas is let in to the mixing cell, the geometry of the cell, and if any active mixing, such as a mixing paddle or a cold bath, is utilized. Once the signal has reached a plateau, the signal at the plateau is taken to be the 118 nm flux for the given mixture. Each phase-matching curve is then constructed from the value at each plateau. Each peak of each phase-matching curve represents the phase-matched mixture for the given xenon pressure. At each peak F_1 has been maximized and \mathcal{P}_ω^3 is constant. Equation 6.6 indicates that as \mathcal{P}_{Xe} increases, and thus $\chi_{eff}^{(3)}$, $\mathcal{P}_{3\omega}^3$ should increase as the square of the xenon pressure. Instead, the 118 nm flux reaches a maximum at 24 Torr of xenon. At the peak of the phase-matching curve for 24 Torr of xenon, it is estimated that there are 10s of nJ of 118 nm light. At higher pressures, instead of increased 118 nm signal, the flux decreases until very little 118 nm light is observed. This same turnover in 118 nm light production versus total pressure was observed in Shi *et al.* [201]. The large discrepancy between the predicted and measured yields of 118 nm light at high pressures indicates that some process is significantly limiting the conversion efficiency. If this process can be understood and counteracted, then signal in molecular detection techniques that use 118 nm could be greatly enhanced.

The decrease in 118 nm flux is accompanied by the observation of a fluorescence streak (Figure 6.18) near the focus of the 355 nm light. The streak follows the propagation axis of the 355 nm light. The streak could be observed and studied through a ConFlat window located above the 355 nm light focal region. The streak became brighter as the 118 nm signal decreases with higher xenon pressures.

A photo diode and an Ocean Optics USB2000 spectrometer were used to examine the fluorescence streak. During phase-matching measurements with xenon pressures greater than 20 torr,




Figure 6.17: Phase-matching curves at different pressures of Xe. While theory indicates that the phase-matched peak heights should increase as the square of the xenon pressure, this is not the case, and 118 nm signal reaches a peak at a phase-matched mixture with 24 torr of xenon.




Figure 6.18: A picture of the fluorescence streak that appears near the 355 nm focal point.

the streak intensity is correlated with 118 nm production. Figure 6.19 shows the fluorescence signal closely following the shape of the phase-matching curve. This indicates that 118 nm photons play a role in the creation of the fluorescence streak. Additionally, when examining a phase-matched mixture (a constant 1:11.5 ratio of Xe:Ar with different total pressures) the intensity of the streak increased with increasing xenon pressures (Figure 6.18). The fluorescence intensity increased linearly at a threshold xenon pressure (~ 20 torr), which is near the peak of the 118 nm photon flux. The intensity of the fluorescence increases linearly as 118 nm photon flux decreases.

In order to further understand the fluorescence streak, we examined 118 nm photon production with a different phase-matching gas, krypton. Krypton is positively dispersive at 118 nm just like argon. It has a measured phase matching ratio of 1:4.6, which agreed with the krypton dispersion calculated in Mahon *et al.* [205]. The same measurements that were conducted with the xenon-argon mixture were also conducted using xenon-krypton. The red curves in Figure 6.20 show the measurements for krypton. This mixture did not yield more 118 nm light than the Xe/Ar mixture and the fluorescence streak was still present. The overall 118 photon flux for phase-matched mixtures reaches a lower peak at a lower xenon pressure than for the xenon-argon mixture. And similar to the xenon-argon case, a fluorescence streak appeared starting with ~ 20 torr of xenon. The xenon-krypton mixture did not have any advantages over the xenon-argon system due to the decrease in the observed 118 nm signal.

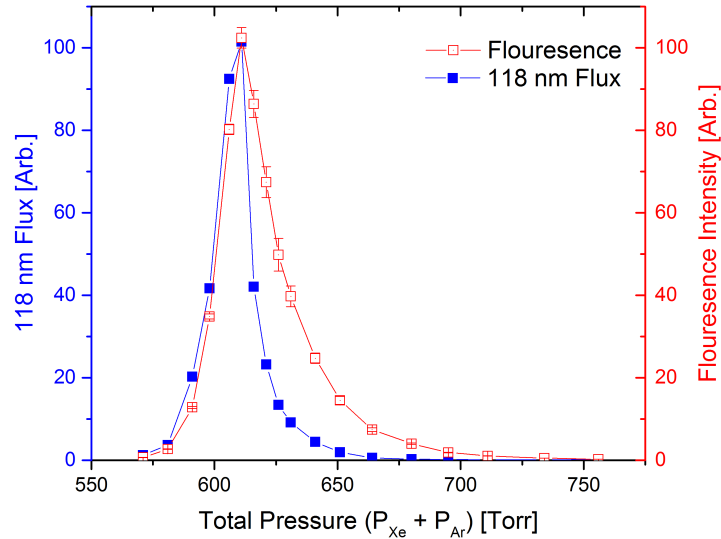


Figure 6.19: A phase-matching curve, with $\mathcal{P}_{Xe} = 50$ torr, taken while monitoring both the 118 nm flux and the fluorescence intensity. The fluorescence signal is correlated with the presence of 118 nm light.

Figures 6.19 and 6.20 suggest that 118 nm light is being created, but is being consumed by some mechanism before it is able to reach the detection region. This is supported by the fact that




Figure 6.20: A comparison of the 118 nm signal (solid markers) and fluorescence streak (open markers) signal at different xenon pressures for a phase-matched mixture with Ar (green, Xe:Ar = 1:11.3) and Kr (red, Xe:Kr = 1:4.6).

the fluorescence streak does not occur when 355 nm light is focused into a cell containing only xenon or only argon. With only one gas, there is the same 355 nm power, \mathcal{P}_ω^3 , but no 118 nm light should be produced due to the lack of a phase-matched mixture. Since no fluorescence streak appeared in this situation, 118 nm light is a prerequisite for the formation of the streak.

The mechanism for this decrease in 118 nm flux above $\mathcal{P}_{Xe} = 20$ torr must include a runaway or positive feedback process since 118 nm signal continues to decrease with increasing fluorescence. If the 118 nm signal did not turnover, but instead reached a plateau, the mechanism would have to be a regulating or a negative feedback process. A plateau indicates that the process requires a threshold amount of 118 nm light to begin, but would stop once it depletes the 118 nm light below the threshold. Instead, we observe a turnover, this indicates that once the mechanism is initiated it “runs away” and consumes as much a 118 nm light as possible. As we increase the xenon pressure for a phase-matched mixture, more 118 nm light should be produced. However, once the fluorescence streak appears, increased pressure leads to a decrease in 118 nm flux.

To gain additional information about the fluorescence streak, a spectrum of the streak was




Figure 6.21: A spectrum of the fluorescence streak created in a xenon-argon mixture. The location of Xe I (Ar I) spectral lines are indicated in green (red).

taken with an Ocean Optics USB2000 spectrometer. The spectral lines from the streak are similar in shape and intensity to those observed in a xenon arc lamp indicating that the fluorescence is most likely from electron-ion recombination in xenon. Figure 6.21 shows the recorded spectrum and the location of the most prominent Xe I and Ar I lines. The spectrum shows predominantly Xe I spectral lines and very weak Ar I spectral lines. Although xenon is present in a much lower concentrations in the gas mixture, it is preferentially ionized over argon. This suggests that a photon above the IP of xenon and below the IP of argon (15.76 eV) [206] is present in the system. 118 nm (10.51 eV) light is not enough energy to ionize xenon, which has an ionization potential (IP) of 12.13 eV [206].

A spectrum was also recorded for xenon-krypton system, Figure 6.22. While there are some Kr I lines in the spectrum, again Xe I lines are the most prominent. Since krypton is not ionized here, this limits the theory for the secondary process in the cell to a xenon only ionizing process, which should fall between the IP of xenon (12.13 eV) and krypton (14.00 eV).




Figure 6.22: A spectrum of the fluorescence streak created in a xenon-krypton mixture. The location of Xe I (Kr I) spectral lines are indicated in green (red).

6.5.4 Possible explanations for the limitations

Despite the data on the fluorescence streak and the 118 nm flux turnover, the mechanism behind this effect is not well understood. Limitations to the conversion efficiency could arise from linear one-photon absorption as well as other non-linear processes not considered in the basic theoretical model discussed in Section 6.5.1. The mechanism for signal decrease must provide an explanation of all observed properties of the fluorescence streak including:

- (1) the predominance of Xe I lines with few Ar I/Kr I lines,
- (2) correlation of the fluorescence streak intensity with the phase matching curve,
- (3) and increased fluorescence with decreased 118 nm production (*ie.* a positive feedback process not a negative feedback process.)

Non-linear processes including the Kerr effect and bandwidth limitations were discussed in [155], but did not sufficiently explain the magnitude and character of the observed decrease. The

spectra in Figures 6.21 and 6.22 suggest that a higher energy photon, possibly from higher-order harmonics, is present. This criterion eliminates the possibility of efficient higher-order harmonic generation (such as fifth-harmonic generation at 17.48 eV) since that photon would ionize both xenon and argon. In that case, more Ar I spectral lines would be expected since argon is more than 11 times more prevalent in the sample. Higher-order harmonic processes are also already unlikely since the gas mixture is not phase matched for these processes. We consider two possible mechanisms, linear absorption of 118 nm light and sum-frequency generation of 88 nm light, that can be tested experimentally.

6.5.4.1 Linear absorption of 118 nm light

Linear reabsorption of 118 nm light by the xenon-argon gas mixture could be an explanation for the decrease in 118 nm flux. Although the 118 nm light is detuned from any transition in xenon or argon, absorption can still be possible in the wings of resonance lines. The most likely transitions for linear absorption of 118 nm in xenon is the $(^2P_{3/2})5d \left[\frac{3}{2} \right]$ line at 119.203 nm. There is an additional $(^2P_{3/2})7s$ line at 117.043 nm, but it is less likely to contribute due to its smaller, by a factor of 4, oscillator strength. The nearest argon line is at 106.66 nm, which is significantly further away and even less likely to contribute to 118 nm light absorption. A more detailed description and calculation of 118 nm linear absorption in a xenon-argon mixture can be found in [155].

To test for linear absorption of 118 nm light, a short absorption cell was placed in between the 118 nm generation cell and the UHV detection chamber. Figure 6.23 shows a schematic of this setup. A MgF window is placed after the MgF lens. The area between the window and lens serves as the absorption cell. By filling this volume with various pressures of xenon and comparing the resulting signal, we can examine the possibility of 118 nm linear absorption by xenon. The 118 nm generation cell, including the sealing MgF lens has been shifted away from the detection region by ~ 2 cm. To account for this change, the 355 nm lens on the outside of the vacuum chamber can be adjusted such that the 118 nm light remains focused at the center of the TOFMS. Additionally, this adjustment slightly changes the Rayleigh range of the focused beam. However, since this test




Figure 6.23: A cartoon of the 118 nm gas mixing cell, absorption cell, and ionization and detection region. The absorption cell is sealed off using a MgF_2 lens and window. The gas cell remains the same length as in Figure 6.16, but has been shifted ~ 2 cm away from the detection chamber.

was performed using NH_3 as the detection ion, the single-photon process is less dependent on the focus of the beam.

Preliminary studies using the absorption cell suggest that linear absorption is unlikely to be the cause of the decrease in 118 nm flux [154]. Adding xenon, the most likely linear absorption candidate, to the absorption cell did not decrease the observed 118 nm flux. Instead the flux increased. The addition of various other gases (argon, nitrogen, air) to the absorption cell also caused an increase in the observed signal. While the increase in signal could be due to a systematic effect, it is unlikely that there is also linear absorption occurring simultaneously.

6.5.4.2 Sum-frequency generation

A second possible explanation for the decrease in 118 nm light production is sum-frequency generation (SFG) from 355 nm and 118 nm photons. This would produce an extreme ultraviolet (EUV) photon at 88 nm (14 eV), which satisfies criteria (1). Although SFG is a $\chi^{(2)}$ processes, which is normally forbidden in atomic gases, second-harmonic and sum-frequency effects have been observed in several atomic gases and metal vapors. These processes were first observed using external electric [207] or magnetic [208, 209] fields to break the symmetry of the gas and

permit second-harmonic generation (SHG). SHG in an atomic gas without an external field was subsequently observed in thallium by Mossberg *et al.* [210]. Additionally, SHG in gases without an external field has also been observed in metal vapors of sodium [211] and zinc [212], and atomic gases of xenon [213,214] and hydrogen [215]. Lastly, one case of Sum-frequency generation in an atomic gas has been observed in sodium [216].

Various processes have been proposed to explain these normally $\chi^{(2)}$ processes in atomic gases. Primarily these processes include symmetry breaking components that allow for normally forbidden process to proceed. The symmetry breaking processes fall into two categories: those that involve photoionization and those with only neutral atoms [217]. The first category includes models where charge separation occurs as a result of multiphoton ionization of the gas. A symmetry breaking electric field is formed due to the separation of the ion and electron. This occurs either from the radiative forces exerted by the pump light or the resulting kinetic energy of the ionization process. The second category includes models that take into account the dipole-forbidden second-order susceptibility of the medium. By including contributions from higher-multipole orders, such as the electric quadrupole and magnetic-dipole of the non-linear medium, $\chi^{(2)}$ maybe non-vanishing [216]. Here, we concentrate on only the first category since Bethune [217] found that only ionization could produce a strong SHG signal.

In previous atomic SHG experiments, it was not always possible to determine which of the exact mechanisms caused SHG in the non-metallic gases. However, there are two conditions that always made SHG more probable [213–215]. The first is ionization of the atomic gas that causes charge separation. The electric field between the ion and electron is sufficient to break the symmetry and allow for SHG. The second condition is a resonant two-photon allowed transition in the atom at the wavelength of the SHG light which helps to enhance the creation of SHG light.

Both of these conditions can be fulfilled in the xenon-argon 118 nm system and could explain the decrease in detected 118 nm flux. In SHG in a single atomic gas, the atom provides both the electric field, usually through multiphoton ionization, and the resonant transition. However, in our case, xenon and argon each provide one of these components, and together aid in the generation

of 88 nm light. The initial electric field may come from 118 nm ionization of contaminants in the gas system that have lower ionization potentials. Once a small field is present one of the $3p^53d \leftarrow 1S_0$ transitions in argon, which are in the 87-89 nm range, can provide the resonant two-photon transition to help enhance SFG. Now that both conditions that make SHG more probable are present, 88 nm light can be created. As soon as 88 nm light is present in the system, it would ionize xenon and increase the symmetry breaking electric field. This mechanism meets criteria (2) and (3). With more 118 nm production there is a higher chance of 88 nm light and thus the fluorescence streak. It also contains a positive feedback loop of which both xenon and argon play a role. The loop is depicted in Fig. 6.24. Although more 118 nm light is created at higher pressures, more EUV light is also created at higher pressures. If the SFG process dominates over the THG process, it can explain the observed fluorescence increase and VUV decrease.

All the measurements with krypton are consistent with this mechanism since there are many more potentially resonant transition in krypton than in argon. Since the SFG photon (14 eV) is right at the ionization potential of krypton (14.00 eV), there could be many high lying excited states of krypton that could assist in SFG.




Figure 6.24: Diagram depicting the feedback mechanism of electric field induced SFG. Symmetry is broken by an electric field caused by ionization of contaminants, which leads to SFG of 88 nm aided by a near resonant transition in argon. The 88 nm light ionizes Xe, which causes a larger electric field and more 88 nm light.

At this point it is uncertain which $3p^53d \leftarrow 1S_0$ transitions in argon would be the one to enhance the SFG process, but the $3p^53d [5/2]_2$ state appears to be the most likely. Using the same argument as in Shi *et al.* [214], the highest lying J=0 or 2 state that also has an efficient one-photon coupling, should have the most efficient SFG process. Although in Shi *et al.*, they examined

second-harmonic generation, the arguments should still be relevant here. For 88 nm light and argon, this would be the $3p^53d[5/2]_2$ state at $113425.964\text{ cm}^{-1}$ (88.165 nm) coupling to $3p^55s[3/2]_1$ at 113643.26 cm^{-1} (87.995 nm), which is 217.3 cm^{-1} away.

A possible experimental test of the SFG theory could be completed with a xenon-neon cell. THG in a phase-matched mixture of xenon-neon has the potential to avoid any resonant transitions with SFG light. The first excited state of neon (16.62 eV) and the IP of neon (21.56 eV) are sufficiently large enough that there are no resonant states to increase the efficiency of SFG production at 88 nm (14 eV). To test a xenon-neon system, measurements would ideally be taken with $\mathcal{P}_{Xe} > 24\text{ torr}$. However, a phase-matched ratio of xenon-neon containing 50 torr of xenon would have a total pressure around 20 bar. This is not achievable in the current experimental setup since conflat components are not designed to hold positive pressures. Instead, a reinforced high pressure cell would have to be constructed, and special considerations would be required for the sealing lens since it would encounter even greater pressure differentials.

One other option for testing the SFG mechanism could be changing the pump laser wavelength. A tunable dye laser around 355 nm could be used instead of a Nd:YAG. If 118 nm light yield is dependent on the pump laser wavelength, specifically if resonances are observed, this would support the SFG mechanism. It would also help to determine which transition in argon is assisting in the SFG process. The pump beam could then be tuned so that the SFG wavelength is sufficiently far away from the resonant 88 nm argon transition. While all the tests to determine whether SFG is the cause of the decrease in 118 nm light signal are experimentally difficult to perform, if completed they would significantly further our understanding of the 118 nm suppression mechanism. If it is possible to overcome the suppression mechanism then many chemical studies that use 118 nm light for ionization and molecule detection would benefit greatly.

Bibliography

- [1] M. H. Anderson, J. R. Ensher, M. R. Matthews, C. E. Wieman, and E. A. Cornell. Observation of bose-einstein condensation in a dilute atomic vapor. Science, 269(5221):198–201, 1995.
- [2] K. B. Davis, M. O. Mewes, M. R. Andrews, N. J. van Druten, D. S. Durfee, D. M. Kurn, and W. Ketterle. Bose-einstein condensation in a gas of sodium atoms. Phys. Rev. Lett., 75:3969–3973, Nov 1995.
- [3] Immanuel Bloch, Jean Dalibard, and Wilhelm Zwerger. Many-body physics with ultracold gases. Rev. Mod. Phys., 80:885–964, Jul 2008.
- [4] M. Takamoto, F. L. Hong, R. Higashi, and H. Katori. An optical lattice clock. Nature, 435:321–324, 2005.
- [5] A. D. Ludlow, T. Zelevinsky, G. K. Campbell, S. Blatt, M. M. Boyd, M. H. G. de Miranda, M. J. Martin, J. W. Thomsen, S. M. Foreman, Jun Ye, T. M. Fortier, J. E. Stalnaker, S. A. Diddams, Y. Le Coq, Z. W. Barber, N. Poli, N. D. Lemke, K. M. Beck, and C. W. Oates. Sr lattice clock at 1×10^{-16} fractional uncertainty by remote optical evaluation with a Ca clock. Science, 319(5871):1805–1808, 2008.
- [6] G. Edward Marti, Ross B. Hutson, Akihisa Goban, Sara L. Campbell, Nicola Poli, and Jun Ye. Imaging optical frequencies with 100 μ Hz precision and 1.1 μ m resolution. Phys. Rev. Lett., 120:103201, Mar 2018.
- [7] B. DeMarco and D. S. Jin. Onset of Fermi degeneracy in a trapped atomic gas. Science, 285(5434):1703–1706, 1999.
- [8] C. A. Regal, M. Greiner, and D. S. Jin. Observation of resonance condensation of fermionic atom pairs. Phys. Rev. Lett., 92:040403, Jan 2004.
- [9] In Rui Rausto, editor, Low Temperature Molecular Spectroscopy, volume 483 of Nato Science Series C. Springer Netherlands, 1993.
- [10] Eric R. Hudson, H. J. Lewandowski, Brian C. Sawyer, and Jun Ye. Cold molecule spectroscopy for constraining the evolution of the fine structure constant. Phys. Rev. Lett., 96:143004, Apr 2006.
- [11] S Truppe, R J Hendricks, S K Tokunaga, H J Lewandowski, M G Kozlov, C Henkel, E A Hinds, and M R Tarbutt. A search for varying fundamental constants using hertz-level frequency measurements of cold CH molecules. Nature Communications, 4(2600), 2013.

- [12] J. Baron, W. C. Campbell, D. DeMille, J. M. Doyle, G. Gabrielse, Y. V. Gurevich, P. W. Hess, N. R. Hutzler, E. Kirilov, I. Kozyryev, B. R. O’Leary, C. D. Panda, M. F. Parsons, E. S. Petrik, B. Spaun, A. C. Vutha, and A. D. West. Order of magnitude smaller limit on the electric dipole moment of the electron. Science, 343(6168):269–272, 2014.
- [13] William B. Cairncross, Daniel N. Gresh, Matt Grau, Kevin C. Cossel, Tanya S. Roussy, Yiqi Ni, Yan Zhou, Jun Ye, and Eric A. Cornell. Precision measurement of the electron’s electric dipole moment using trapped molecular ions. Phys. Rev. Lett., 119:153001, Oct 2017.
- [14] Lotte Holmegaard, Jens H. Nielsen, Iftach Nevo, Henrik Stapelfeldt, Frank Filsinger, Jochen Küpper, and Gerard Meijer. Laser-induced alignment and orientation of quantum-state-selected large molecules. Phys. Rev. Lett., 102:023001, Jan 2009.
- [15] Frank Filsinger, Jochen Kpper, Gerard Meijer, Lotte Holmegaard, Jens H. Nielsen, Iftach Nevo, Jonas L. Hansen, and Henrik Stapelfeldt. Quantum-state selection, alignment, and orientation of large molecules using static electric and laser fields. The Journal of Chemical Physics, 131(6):064309, 2009.
- [16] Piotr S. Żuchowski and Jeremy M. Hutson. Low-energy collisions of NH_3 and ND_3 with ultracold Rb atoms. Phys. Rev. A, 79:062708, Jun 2009.
- [17] M. H. G. de Miranda, A. Chotia, B. Neyenhuis, D. Wang, G. Qumner, S. Ospelkaus, J. L. Bohn, J. Ye, and D. S. Jin. Controlling the quantum stereodynamics of ultracold bimolecular reactions. Nature Physics, 7, 2011.
- [18] J. Doyle, B. Friedrich, R. V. Krems, and F. Masnou-Seeuws. Editorial: Quo vadis, cold molecules? The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics, 31(2):149–164, Nov 2004.
- [19] Lincoln D Carr, David DeMille, Roman V Krems, and Jun Ye. Cold and ultracold molecules: science, technology and applications. New Journal of Physics, 11(5):055049, 2009.
- [20] Melanie Schnell and Gerard Meijer. Cold molecules: Preparation, applications, and challenges. Angewandte Chemie International Edition, 48(33):6010–6031, 2009.
- [21] Roman Krems, Bretislav Friedrich, and William C Stwalley. Cold Molecules: Theory, Experiment, Applications. CRC Press, 2009.
- [22] Stephen D. Hogan, Michael Motsch, and Frederic Merkt. Deceleration of supersonic beams using inhomogeneous electric and magnetic fields. Phys. Chem. Chem. Phys., 13:18705–18723, 2011.
- [23] Deborah S. Jin and Jun Ye. Introduction to ultracold molecules: New frontiers in quantum and chemical physics. Chemical Reviews, 112(9):4801–4802, 2012.
- [24] Sebastiaan Y. T. van de Meerakker, Hendrick L. Bethlem, Nicolas Vanhaecke, and Gerard Meijer. Manipulation and control of molecular beams. Chemical Reviews, 112(9):4828–4878, 2012.
- [25] Justin Jankunas and Andreas Osterwalder. Cold and controlled molecular beams: Production and applications. Annual Review of Physical Chemistry, 66(1):241–262, 2015.

- [26] T E Wall. Preparation of cold molecules for high-precision measurements. Journal of Physics B: Atomic, Molecular and Optical Physics, 49(24):243001, 2016.
- [27] John L. Bohn, Ana Maria Rey, and Jun Ye. Cold molecules: Progress in quantum engineering of chemistry and quantum matter. Science, 357(6355):1002–1010, 2017.
- [28] William C. Stwalley and He Wang. Photoassociation of ultracold atoms: A new spectroscopic technique. Journal of Molecular Spectroscopy, 195(2):194 – 228, 1999.
- [29] K-K Ni, S Ospelkaus, M. H. G. de Miranda, A Pe’er, B. Neyenhuis, J. J. Zirbel, S. Kotochigova, P. S. Julienne, D. S. Jin, and J. Ye. A high phase-space-density gas of polar molecules. Science, 322:231–235, 2008.
- [30] Thorsten Köhler, Krzysztof Góral, and Paul S. Julienne. Production of cold molecules via magnetically tunable Feshbach resonances. Rev. Mod. Phys., 78:1311–1361, Dec 2006.
- [31] K.-K. Ni, S. Ospelkaus, D. Wang, G. Qummer, B. Neyenhuis, M. H. G. de Miranda, J. L. Bohn, J. Ye, and D. S. Jin. Dipolar collisions of polar molecules in the quantum regime. Nature, 464:1324–1328, 2010.
- [32] Alejandra L Collopy, Matthew T Hummon, Mark Yeo, Bo Yan, and Jun Ye. Prospects for a narrow line MOT in YO. New Journal of Physics, 17(5):055008, 2015.
- [33] J. Lim, J. R. Almond, M. A. Trigatzis, J. A. Devlin, N. J. Fitch, B. E. Sauer, M. R. Tarbutt, and E. A. Hinds. Laser cooled ybf molecules for measuring the electron’s electric dipole moment. Phys. Rev. Lett., 120:123201, Mar 2018.
- [34] Benjamin K. Stuhl, Mark Yeo, Brian C. Sawyer, Matthew T. Hummon, and Jun Ye. Microwave state transfer and adiabatic dynamics of magnetically trapped polar molecules. Phys. Rev. A, 85:033427, Mar 2012.
- [35] M. Zeppenfeld, M. Motsch, P. W. H. Pinkse, and G. Rempe. Optoelectrical cooling of polar molecules. Phys. Rev. A, 80:041401, Oct 2009.
- [36] Dima Egorov, Thierry Lahaye, Wieland Schöllkopf, Bretislav Friedrich, and John M. Doyle. Buffer-gas cooling of atomic and molecular beams. Phys. Rev. A, 66:043401, Oct 2002.
- [37] David Patterson and John M. Doyle. Bright, guided molecular beam with hydrodynamic enhancement. The Journal of Chemical Physics, 126(15):154307, 2007.
- [38] Nicholas R. Hutzler, Hsin-I Lu, and John M. Doyle. The buffer gas beam: An intense, cold, and slow source for atoms and molecules. Chemical Reviews, 112(9):4803–4827, 2012.
- [39] E. S. Shuman, J. F. Barry, and D. DeMille. Laser cooling of a diatomic molecule. Nature, 467:820–823, 2010.
- [40] Matthew T. Hummon, Mark Yeo, Benjamin K. Stuhl, Alejandra L. Collopy, Yong Xia, and Jun Ye. 2D magneto-optical trapping of diatomic molecules. Phys. Rev. Lett., 110:143001, Apr 2013.
- [41] V. Zhelyazkova, A. Cournol, T. E. Wall, A. Matsushima, J. J. Hudson, E. A. Hinds, M. R. Tarbutt, and B. E. Sauer. Laser cooling and slowing of CaF molecules. Phys. Rev. A, 89:053416, May 2014.

- [42] J. F. Barry, D. J. McCarron, E. B. Norrgard, M. H. Steinecker, and D. DeMille. Magneto-optical trapping of a diatomic molecule. Nature, 512:286–289, 2014.
- [43] M. I. Fabrikant, Tian Li, N. J. Fitch, N. Farrow, Jonathan D. Weinstein, and H. J. Lewandowski. Method for traveling-wave deceleration of buffer-gas beams of CH. Phys. Rev. A, 90:033418, Sep 2014.
- [44] Piergiorgio Casavecchia. Chemical reaction dynamics with molecular beams. Reports on Progress in Physics, 63(3):355, 2000.
- [45] A. W. Kleyn. Molecular beams and chemical dynamics at surfaces. Chem. Soc. Rev., 32:87–95, 2003.
- [46] Timothy J. McCarthy, Michael T. Timko, and Dudley R. Herschbach. Manipulation of slow molecular beams by static external fields. The Journal of Chemical Physics, 125(13):133501, 2006.
- [47] Hendrick L. Bethlem, André J. A. van Roij, Rienk T. Jongma, and Gerard Meijer. Alternate gradient focusing and deceleration of a molecular beam. Phys. Rev. Lett., 88:133003, Mar 2002.
- [48] M. R. Tarbutt, H. L. Bethlem, J. J. Hudson, V. L. Ryabov, V. A. Ryzhov, B. E. Sauer, G. Meijer, and E. A. Hinds. Slowing heavy, ground-state molecules using an alternating gradient decelerator. Phys. Rev. Lett., 92:173002, Apr 2004.
- [49] Hendrick L. Bethlem, Jacqueline van Veldhoven, Melanie Schnell, and Gerard Meijer. Trapping polar molecules in an ac trap. Phys. Rev. A, 74:063403, Dec 2006.
- [50] T. E. Wall, J. F. Kanem, J. M. Dyne, J. J. Hudson, B. E. Sauer, E. A. Hinds, and M. R. Tarbutt. Stark deceleration of CaF molecules in strong- and weak-field seeking states. Phys. Chem. Chem. Phys., 13:18991–18999, 2011.
- [51] Kirstin Wohlfart, Frank Filsinger, Fabian Grätz, Jochen Küpper, and Gerard Meijer. Stark deceleration of OH radicals in low-field-seeking and high-field-seeking quantum states. Phys. Rev. A, 78:033421, Sep 2008.
- [52] Kirstin Wohlfart, Fabian Grätz, Frank Filsinger, Henrik Haak, Gerard Meijer, and Jochen Küpper. Alternating-gradient focusing and deceleration of large molecules. Phys. Rev. A, 77:031404, Mar 2008.
- [53] S.R Procter, Y Yamakita, F Merkt, and T.P Softley. Controlling the motion of hydrogen molecules. Chemical Physics Letters, 374(5):667 – 675, 2003.
- [54] E. Vliegen, H. J. Wörner, T. P. Softley, and F. Merkt. Nonhydrogenic effects in the deceleration of Rydberg atoms in inhomogeneous electric fields. Phys. Rev. Lett., 92:033005, Jan 2004.
- [55] E Vliegen and F Merkt. Stark deceleration of hydrogen atoms. Journal of Physics B: Atomic, Molecular and Optical Physics, 39(11):L241, 2006.
- [56] E. Vliegen and F. Merkt. Normal-incidence electrostatic Rydberg atom mirror. Phys. Rev. Lett., 97:033002, Jul 2006.

- [57] Nicolas Vanhaecke, Urban Meier, Markus Andrist, Beat H. Meier, and Frédéric Merkt. Multistage zeeman deceleration of hydrogen atoms. Phys. Rev. A, 75:031402, Mar 2007.
- [58] S D Hogan, A W Wiederkehr, M Andrist, H Schmutz, and F Merkt. Slow beams of atomic hydrogen by multistage zeeman deceleration. Journal of Physics B: Atomic, Molecular and Optical Physics, 41(8):081005, 2008.
- [59] Edvardas Narevicius, Adam Libson, Christian G. Parthey, Isaac Chavez, Julia Narevicius, Uzi Even, and Mark G. Raizen. Stopping supersonic beams with a series of pulsed electromagnetic coils: An atomic coilgun. Phys. Rev. Lett., 100:093003, Mar 2008.
- [60] S. A. Rangwala, T. Junglen, T. Rieger, P. W. H. Pinkse, and G. Rempe. Continuous source of translationally cold dipolar molecules. Phys. Rev. A, 67:043406, Apr 2003.
- [61] T. Junglen, T. Rieger, S. A. Rangwala, P. W. H. Pinkse, and G. Rempe. Slow ammonia molecules in an electrostatic quadrupole guide. The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics, 31(2):365–373, Nov 2004.
- [62] T. Rieger, T. Junglen, S. A. Rangwala, P. W. H. Pinkse, and G. Rempe. Continuous loading of an electrostatic trap for polar molecules. Phys. Rev. Lett., 95:173002, Oct 2005.
- [63] Jeffrey J. Kay, Sebastiaan Y. T. van de Meerakker, Kevin E. Strecker, and David W. Chandler. Production of cold ND₃ by kinematic cooling. Faraday Discuss., 142:143–153, 2009.
- [64] Michael S. Elioff, James J. Valentini, and David W. Chandler. Subkelvin cooling NO molecules via "billiard-like" collisions with argon. Science, 302(5652):1940–1943, 2003.
- [65] Alexandre Trottier, David Carty, and Eckart Wrede. Photostop: production of zero-velocity molecules by photodissociation in a molecular beam. Molecular Physics, 109(5):725–733, 2011.
- [66] P. F. Barker and M. N. Shneider. Slowing molecules by optical microlinear deceleration. Phys. Rev. A, 66:065402, Dec 2002.
- [67] R. Fulton, A. I. Bishop, and P. F. Barker. Optical Stark decelerator for molecules. Phys. Rev. Lett., 93:243004, Dec 2004.
- [68] A I Bishop, L Wang, and P F Barker. Creating cold stationary molecular gases by optical Stark deceleration. New Journal of Physics, 12(7):073028, 2010.
- [69] Yaling Yin, Qi Zhou, Lianzhong Deng, Yong Xia, and Jianping Yin. Multistage optical Stark decelerator for a pulsed supersonic beam with a quasi-cw optical lattice. Opt. Express, 17(13):10706–10717, Jun 2009.
- [70] Lianzhong Deng, Shunyong Hou, and Jianping Yin. Optical Stark decelerator for molecules with a traveling potential well. Phys. Rev. A, 95:033409, Mar 2017.
- [71] Manish Gupta and Dudley Herschbach. Slowing and speeding molecular beams by means of a rapidly rotating source. The Journal of Physical Chemistry A, 105(9):1626–1637, 2001.
- [72] M. Strebels, F. Stienkemeier, and M. Mudrich. Improved setup for producing slow beams of cold molecules using a rotating nozzle. Phys. Rev. A, 81:033409, Mar 2010.

- [73] S. Chervenkov, X. Wu, J. Bayerl, A. Rohlfes, T. Gantner, M. Zeppenfeld, and G. Rempe. Continuous centrifuge decelerator for polar molecules. Phys. Rev. Lett., 112:013001, Jan 2014.
- [74] Xing Wu, Thomas Gantner, Manuel Koller, Martin Zeppenfeld, Sotir Chervenkov, and Gerhard Rempe. A cryofuge for cold-collision experiments with slow polar molecules. Science, 358(6363):645–648, 2017.
- [75] Joop J. Gilijamse, Steven Hoekstra, Sebastiaan Y. T. van de Meerakker, Gerrit C. Groenenboom, and Gerard Meijer. Near-threshold inelastic collisions using molecular beams with a tunable velocity. Science, 313(5793):1617–1620, 2006.
- [76] Sebastiaan Y T van de Meerakker, Irena Labazan, Steven Hoekstra, Jochen Kpper, and Gerard Meijer. Production and deceleration of a pulsed beam of metastable NH ($a^1\Delta$) radicals. Journal of Physics B: Atomic, Molecular and Optical Physics, 39(19):S1077, 2006.
- [77] Ludwig Scharfenberg, Sebastiaan Y. T. van de Meerakker, and Gerard Meijer. Crossed beam scattering experiments with optimized energy resolution. Phys. Chem. Chem. Phys., 13:8448–8456, 2011.
- [78] L. Scharfenberg, K. B. Gubbels, M. Kirste, G. C. Groenenboom, A. van der Avoird, G. Meijer, and S. Y. T. van de Meerakker. Scattering of Stark-decelerated oh radicals with rare-gas atoms. Eur. Phys. J. D, 65(1):189–198, Nov 2011.
- [79] Simon Merz, Nicolas Vanhaecke, Wolfgang Jäger, Melanie Schnell, and Gerard Meijer. Decelerating molecules with microwave fields. Phys. Rev. A, 85:063411, Jun 2012.
- [80] Floris M. H. Cromptoets, Hendrick L. Bethlem, Rienk T. Jongma, and Gerard Meijer. A prototype storage ring for neutral molecules. Nature, 411:174–176, May 2001.
- [81] Peter C. Zieger, Sebastiaan Y. T. van de Meerakker, Cynthia E. Heiner, Hendrick L. Bethlem, André J. A. van Roij, and Gerard Meijer. Multiple packets of neutral molecules revolving for over a mile. Phys. Rev. Lett., 105:173001, Oct 2010.
- [82] Cynthia E. Heiner, David Carty, Gerard Meijer, and Hendrick L. Bethlem. A molecular synchrotron. Nat Phys, 3:115118, 2007.
- [83] Shunyong Hou, Bin Wei, Lianzhong Deng, and Jianping Yin. A novel molecular synchrotron for cold collisions and edm experiments. Sci Rep, 6:32663, 2016.
- [84] Hendrick L. Bethlem, Giel Berden, Floris M. H. Cromptoets, Rienk T. Jongma, Andre J. A. van Roij, and Gerard Meijer. Electrostatic trapping of ammonia molecules. Nature, 406, Aug 2000.
- [85] Jacqueline van Veldhoven, Hendrick L. Bethlem, Melanie Schnell, and Gerard Meijer. Versatile electrostatic trap. Phys. Rev. A, 73:063408, Jun 2006.
- [86] J. J. Gilijamse, S. Hoekstra, N. Vanhaecke, S.Y.T. van de Meerakker, and G. Meijer. Loading Stark-decelerated molecules into electrostatic quadrupole traps. Eur. Phys. J. D, 57(1):33–41, Mar 2010.

- [87] Brian C. Sawyer, Benjamin L. Lev, Eric R. Hudson, Benjamin K. Stuhl, Manuel Lara, John L. Bohn, and Jun Ye. Magneto-electrostatic trapping of ground state OH molecules. *Phys. Rev. Lett.*, 98:253002, Jun 2007.
- [88] J. Riedel, S. Hoekstra, W. Jäger, J. J. Gilijamse, S. Y. T. van de Meerakker, and G. Meijer. Accumulation of Stark-decelerated nh molecules in a magnetic trap. *Eur. Phys. J. D*, 65(1):161–166, Nov 2011.
- [89] Jacqueline van Veldhoven, Hendrick L. Bethlem, and Gerard Meijer. ac electric trap for ground-state molecules. *Phys. Rev. Lett.*, 94:083001, Mar 2005.
- [90] Hendrick L Bethlem, M R Tarbutt, Jochen Kpper, David Carty, Kirstin Wohlfart, E A Hinds, and Gerard Meijer. Alternating gradient focusing and deceleration of polar molecules. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 39(16):R263, 2006.
- [91] Melanie Schnell, Peter Lützow, Jacqueline van Veldhoven, Hendrick L. Bethlem, Jochen Kpper, Bretislav Friedrich, Monika Schleier-Smith, Henrik Haak, and Gerard Meijer. A linear ac trap for polar molecules in their ground state. *J Phys. Chem. A*, 111(31):7411–7419, 2007.
- [92] Brian C. Sawyer, Benjamin K. Stuhl, Dajun Wang, Mark Yeo, and Jun Ye. Molecular beam collisions with a magnetically trapped target. *Phys. Rev. Lett.*, 101:203203, Nov 2008.
- [93] Brian C. Sawyer, Benjamin K. Stuhl, Mark Yeo, Timur V. Tscherebul, Matthew T. Hummon, Yong Xia, Jacek Klos, David Patterson, John M. Doyle, and Jun Ye. Cold heteromolecular dipolar collisions. *Phys. Chem. Chem. Phys.*, 13:19059–19066, 2011.
- [94] L. P. Parazzoli, N. J. Fitch, P. S. Żuchowski, J. M. Hutson, and H. J. Lewandowski. Large effects of electric fields on atom-molecule collisions at millikelvin temperatures. *Phys. Rev. Lett.*, 106:193201, May 2011.
- [95] N. J. Fitch. Traveling-wave Stark-decelerated molecular beams for cold collision experiments. PhD thesis, University of Colorado Boulder, 2013.
- [96] Sebastiaan Y. T. van de Meerakker and Gerard Meijer. Collision experiments with Stark-decelerated beams. *Faraday Discuss.*, 142:113–126, 2009.
- [97] Ludwig Scharfenberg, Jacek Klos, Paul J. Dagdigian, Millard H. Alexander, Gerard Meijer, and Sebastiaan Y. T. van de Meerakker. State-to-state inelastic scattering of Stark-decelerated OH radicals with Ar atoms. *Phys. Chem. Chem. Phys.*, 12:10660–10670, 2010.
- [98] Cunfeng Cheng, Aernout P. P. van der Poel, Paul Jansen, Marina Quintero-Pérez, Thomas E. Wall, Wim Ubachs, and Hendrick L. Bethlem. Molecular fountain. *Phys. Rev. Lett.*, 117:253201, Dec 2016.
- [99] J. van Veldhoven, J. Küpper, H. L. Bethlem, B. Sartakov, A. J. A. van Roij, and G. Meijer. Decelerated molecular beams for high-resolution spectroscopy. *Eur. Phys. J. D*, 31(2):337–349, Nov 2004.
- [100] Moritz Kirste, Xingan Wang, H. Christian Schewe, Gerard Meijer, Kopin Liu, Ad van der Avoird, Liesbeth M. C. Janssen, Koos B. Gubbels, Gerrit C. Groenenboom, and Sebastiaan Y. T. van de Meerakker. Quantum-state resolved bimolecular collisions of velocity-controlled OH with NO radicals. *Science*, 338(6110):1060–1063, 2012.

- [101] Sebastiaan Y. T. van de Meerakker, Nicolas Vanhaecke, Mark P. J. van der Loo, Gerrit C. Groenenboom, and Gerard Meijer. Direct measurement of the radiative lifetime of vibrationally excited OH radicals. Phys. Rev. Lett., 95:013003, Jul 2005.
- [102] Joop J. Gilijamse, Steven Hoekstra, Samuel A. Meek, Markus Metsl, Sebastiaan Y. T. van de Meerakker, Gerard Meijer, and Gerrit C. Groenenboom. The radiative lifetime of metastable CO ($a^3\Pi$, $v=0$). J. Chem. Phys., 127(22):221102, 2007.
- [103] Steven Hoekstra, Joop J. Gilijamse, Boris Sartakov, Nicolas Vanhaecke, Ludwig Scharfenberg, Sebastiaan Y. T. van de Meerakker, and Gerard Meijer. Optical pumping of trapped neutral molecules by blackbody radiation. Phys. Rev. Lett., 98:133001, Mar 2007.
- [104] H. L. Bethlem, M. Kajita, B. Sartakov, G. Meijer, and W. Ubachs. Prospects for precision measurements on ammonia molecules in a fountain. Eur. Phys. J. Spec. Top., 163(1):55–69, Oct 2008.
- [105] M. R. Tarbutt, J. J. Hudson, B. E. Sauer, and E. A. Hinds. Prospects for measuring the electric dipole moment of the electron using electrically trapped polar molecules. Faraday Discuss., 142:37–56, 2009.
- [106] N.J. Fitch, D.A. Esteves, M.I. Fabrikant, T.C. Briles, Y. Shyur, L.P. Parazzoli, and H.J. Lewandowski. State purity of decelerated molecular beams. Journal of Molecular Spectroscopy, 278:1 – 6, 2012.
- [107] Hendrick L. Bethlem, Giel Berden, and Gerard Meijer. Decelerating neutral dipolar molecules. Phys. Rev. Lett., 83:1558–1561, Aug 1999.
- [108] J. R. Bochinski, Eric R. Hudson, H. J. Lewandowski, Gerard Meijer, and Jun Ye. Phase space manipulation of cold free radical OH molecules. Phys. Rev. Lett., 91:243001, Dec 2003.
- [109] Eric R. Hudson, Christopher Ticknor, Brian C. Sawyer, Craig A. Taatjes, H. J. Lewandowski, J. R. Bochinski, J. L. Bohn, and Jun Ye. Production of cold formaldehyde molecules for study and control of chemical reaction dynamics with hydroxyl radicals. Phys. Rev. A, 73:063404, Jun 2006.
- [110] S K Tokunaga, J M Dyne, E A Hinds, and M R Tarbutt. Stark deceleration of lithium hydride molecules. New Journal of Physics, 11(5):055038, 2009.
- [111] Sebastian Jung, Eberhard Tiemann, and Christian Lisdat. Cold atoms and molecules from fragmentation of decelerated SO₂. Phys. Rev. A, 74:040701, Oct 2006.
- [112] Bucicov, O., Nowak, M., Jung, S., Meijer, G., Tiemann, E., and Lisdat, C. Cold SO₂ molecules by Stark deceleration. Eur. Phys. J. D, 46(3):463–469, 2008.
- [113] Xingan Wang, Moritz Kirste, Gerard Meijer, and Sebastiaan Y. T. van de Meerakker. Stark deceleration of NO radicals. Z Phys Chem, 227:15951604, 2013.
- [114] Hendrick L. Bethlem, Floris M. H. Cromptvoets, Rienk T. Jongma, Sebastiaan Y. T. van de Meerakker, and Gerard Meijer. Deceleration and trapping of ammonia using time-varying electric fields. Phys. Rev. A, 65:053416, May 2002.

- [115] Steven Hoekstra, Markus Metsälä, Peter C. Zieger, Ludwig Scharfenberg, Joop J. Gilijamse, Gerard Meijer, and Sebastiaan Y. T. van de Meerakker. Electrostatic trapping of metastable NH molecules. *Phys. Rev. A*, 76:063408, Dec 2007.
- [116] Dominik Haas, Sebastian Scherb, Dongdong Zhang, and Stefan Willitsch. Optimizing the density of Stark decelerated radicals at low final velocities: a tutorial review. *EPJ Techniques and Instrumentation*, 4(1):6, Sep 2017.
- [117] L P Parazzoli, N Fitch, D S Lobser, and H J Lewandowski. High-energy-resolution molecular beams for cold collision studies. *New J. Phys.*, 11(5):055031, 2009.
- [118] Ludwig Scharfenberg, Henrik Haak, Gerard Meijer, and Sebastiaan Y. T. van de Meerakker. Operation of a Stark decelerator with optimum acceptance. *Phys. Rev. A*, 79:023410, Feb 2009.
- [119] Shunyong Hou, Shengqiang Li, Lianzhong Deng, and Jianping Yin. Dependences of slowing results on both decelerator parameters and the new operating mode: taking ND₃ molecules as an example. *J. Phys. B*, 46(4):045301, 2013.
- [120] Dongdong Zhang, Gerard Meijer, and Nicolas Vanhaecke. Advanced switching schemes in a Stark decelerator. *Phys. Rev. A*, 93:023408, Feb 2016.
- [121] Andreas Osterwalder, Samuel A. Meek, Georg Hammer, Henrik Haak, and Gerard Meijer. Deceleration of neutral molecules in macroscopic traveling traps. *Phys. Rev. A*, 81:051401, May 2010.
- [122] Samuel A. Meek, Maxwell F. Parsons, Georg Heyne, Viktor Platschkowski, Henrik Haak, Gerard Meijer, and Andreas Osterwalder. A traveling wave decelerator for neutral polar molecules. *Rev. Sci. Instrum.*, 82(9):093108, 2011.
- [123] J.E. van den Berg, S.C. Mathavan, C. Meinema, J. Nauta, T.H. Nijbroek, K. Jungmann, H.L. Bethlem, and S. Hoekstra. Traveling-wave deceleration of SrF molecules. *J. Mol. Spectrosc.*, 300:22 – 25, 2014.
- [124] J. E. van den Berg, S. H. Turkesteen, E. B. Prinsen, and S. Hoekstra. Deceleration and trapping of heavy diatomic molecules using a ring-decelerator. *Eur. Phys. J. D*, 66(9):235, Sep 2012.
- [125] N. E. Bulleid, R. J. Hendricks, E. A. Hinds, Samuel A. Meek, Gerard Meijer, Andreas Osterwalder, and M. R. Tarbutt. Traveling-wave deceleration of heavy polar molecules in low-field-seeking states. *Phys. Rev. A*, 86:021404, Aug 2012.
- [126] Sreekanth C. Mathavan, Artem Zapara, Quinten Esajas, and Steven Hoekstra. Deceleration of a supersonic beam of SrF molecules to 120ms⁻¹. *ChemPhysChem*, 17(22):3709–3713, 2016.
- [127] Marina Quintero-Pérez, Paul Jansen, Thomas E. Wall, Joost E. van den Berg, Steven Hoekstra, and Hendrick L. Bethlem. Static trapping of polar molecules in a traveling wave decelerator. *Phys. Rev. Lett.*, 110:133003, Mar 2013.
- [128] Paul Jansen, Marina Quintero-Pérez, Thomas E. Wall, Joost E. van den Berg, Steven Hoekstra, and Hendrick L. Bethlem. Deceleration and trapping of ammonia molecules in a traveling-wave decelerator. *Phys. Rev. A*, 88:043424, Oct 2013.

- [129] Marina Quintero-Prez, Thomas E. Wall, Steven Hoekstra, and Hendrick L. Bethlem. Preparation of an ultra-cold sample of ammonia molecules for precision measurements. J Mol. Spectrosc., 300:112 – 115, 2014.
- [130] Congsen Meng, Aernout P. P. van der Poel, Cunfeng Cheng, and Hendrick L. Bethlem. Femtosecond laser detection of Stark-decelerated and trapped methylfluoride molecules. Phys. Rev. A, 92:023404, Aug 2015.
- [131] Qin Wang, Shunyong Hou, Liang Xu, and Jianping Yin. Slowing and cooling of heavy or light (even with a tiny electric dipole moment) polar molecules using a novel, versatile electrostatic Stark decelerator. Phys. Chem. Chem. Phys., 18:5432–5443, 2016.
- [132] Samuel A. Meek, Hendrick L. Bethlem, Horst Conrad, and Gerard Meijer. Trapping molecules on a chip in traveling potential wells. Phys. Rev. Lett., 100:153003, Apr 2008.
- [133] Samuel A. Meek, Horst Conrad, and Gerard Meijer. A Stark decelerator on a chip. New J. Phys., 11(5):055024, 2009.
- [134] Samuel A. Meek, Horst Conrad, and Gerard Meijer. Trapping molecules on a chip. Science, 324(5935):1699–1702, 2009.
- [135] Samuel A. Meek, Gabriele Santambrogio, Boris G. Sartakov, Horst Conrad, and Gerard Meijer. Suppression of nonadiabatic losses of molecules from chip-based microtraps. Phys. Rev. A, 83:033413, Mar 2011.
- [136] S. Marx, D. Adu Smith, G. Inero, S. A. Meek, B. G. Sartakov, G. Meijer, and G. Santambrogio. Measuring and manipulating the temperature of cold molecules trapped on a chip. Phys. Rev. A, 92:063408, Dec 2015.
- [137] Floris M. H. Crompvoets, Rienk T. Jongma, Hendrick L. Bethlem, André J. A. van Roij, and Gerard Meijer. Longitudinal focusing and cooling of a molecular beam. Phys. Rev. Lett., 89:093004, Aug 2002.
- [138] Sean D. S. Gordon and Andreas Osterwalder. 3D-printed beam splitter for polar neutral molecules. Phys. Rev. Applied, 7:044022, Apr 2017.
- [139] Sebastiaan Y. T. van de Meerakker, Nicolas Vanhaecke, Hendrick L. Bethlem, and Gerard Meijer. Transverse stability in a Stark decelerator. Phys. Rev. A, 73:023401, Feb 2006.
- [140] B. C. Sawyer, B. K. Stuhl, B. L. Lev, J. Ye, and E. R. Hudson. Mitigation of loss within a molecular Stark decelerator. The European Physical Journal D, 48(2):197–209, Jul 2008.
- [141] Shunyong Hou, Qin Wang, Lianzhong Deng, and Jianping Yin. A switched ring Stark decelerator for both light and heavy polar molecules. J. Phys. B, 49(6):065301, 2016.
- [142] Hendrick L. Bethlem, Giel Berden, André J. A. van Roij, Floris M. H. Crompvoets, and Gerard Meijer. Trapping neutral molecules in a traveling potential well. Phys. Rev. Lett., 84:5744–5747, Jun 2000.
- [143] Hao Wu, David Reens, Tim Langen, Yuval Shagam, Daniela Fontecha, and Jun Ye. Enhancing radical molecular beams by skimmer cooling. Phys. Chem. Chem. Phys., 20:11615–11621, 2018.

- [144] W. Ronald Gentry and Clayton F. Giese. Tenmicrosecond pulsed molecular beam source and a fast ionization detector. Review of Scientific Instruments, 49(5):595–600, 1978.
- [145] R. L. Byer and M. D. Duncan. A 100 sec, reliable, 10 hz pulsed supersonic molecular beam source. The Journal of Chemical Physics, 74(4):2174–2179, 1981.
- [146] B. Yan, P. F. H. Claus, B. G. M. van Oorschot, L. Gerritsen, A. T. J. B. Eppink, S. Y. T. van de Meerakker, and D. H. Parker. A new high intensity and short-pulse molecular beam valve. Review of Scientific Instruments, 84(2):023102, 2013.
- [147] Charles E. Otis and Philip M. Johnson. A simple pulsed valve for use in supersonic nozzle experiments. Review of Scientific Instruments, 51(8):1128–1129, 1980.
- [148] U. Even. The Even-Lavie valve as a source for high intensity supersonic beam. EPJ Techniques and Instrumentation, 2(1):17, Dec 2015.
- [149] Jon B. Cross and James J. Valentini. High repetition rate pulsed nozzle beam source. Review of Scientific Instruments, 53(1):38–42, 1982.
- [150] D. Proch and T. Trickl. A highintensity multipurpose piezoelectric pulsed molecular beam source. Review of Scientific Instruments, 60(4):713–716, 1989.
- [151] Wolfgang Christen and Klaus Rademann. Cooling and slowing in high-pressure jet expansions. Phys. Rev. A, 77:012702, Jan 2008.
- [152] Joseph M. Beames, Fang Liu, and Marsha I. Lester. 1+1 resonant multiphoton ionisation of OH radicals via the $A^2\Sigma^+$ state: insights from direct comparison with A-X laser-induced fluorescence detection. Molecular Physics, 112(7):897–903, 2014.
- [153] John M. Gray, Jason A. Bossert, Yomay Shyur, and H. J. Lewandowski. Measurements of trap dynamics of cold oh molecules using resonance-enhanced multiphoton ionization. Phys. Rev. A, 96:023416, Aug 2017.
- [154] Ben Saarel. Creating a better source of 118 nm light. Undergraduate honors thesis, University of Colorado Boulder.
- [155] Travis C. Briles. Production, deceleration, and detection of OH radicals. PhD thesis, University of Colorado Boulder.
- [156] Peter Atkins and Ronal Friedman. Molecular Quantum Mechanics. Oxford University Press, 5 edition, 2011.
- [157] A. C. Cheung, D. M. Rank, C. H. Townes, D. D. Thornton, and W. J. Welch. Detection of NH_3 molecules in the interstellar medium by their microwave emission. Phys. Rev. Lett., 21:1701–1705, Dec 1968.
- [158] K.R. Lang and R.F. Willson. NGC 2264 molecular cloud: Ammonia and OH observations. Astrophys. J.; (United States).
- [159] Paul T. P. Ho and Charles H. Townes. Interstellar ammonia. Annual Review of Astronomy and Astrophysics, 21(1):239–270, 1983.

- [160] S. Maret, A. Faure, E. Scifoni, and L. Wiesenfeld. On the robustness of the ammonia thermometer. Monthly Notices of the Royal Astronomical Society, 399(1):425–431, 2009.
- [161] R.L. Bhattacharjee, L.H. Johnston, G.R. Sudhakaran, and J.C. Sarker. Submillimeter laser Stark spectroscopy of $^{14}\text{ND}_3$. Journal of Molecular Spectroscopy, 138(1):38 – 48, 1989.
- [162] C. H. Townes and A. L. Schawlow. Microwave Spectroscopy. Dover Books, 2012.
- [163] Bruce D. Kay and T. D. Raymond. Rotationally inelastic gassurface scattering: NH_3 from NH_3 -saturated $\text{W}(100)$. The Journal of Chemical Physics, 85(7):4140–4147, 1986.
- [164] M.N.R. Ashfold, R.N. Dixon, R.J. Stickland, and C.M. Western. 2+1 MPI spectroscopy of $\tilde{B}^1\text{E}''$ state NH_3 and ND_3 : rotational analysis of the origin bands. Chemical Physics Letters, 138(2):201 – 208, 1987.
- [165] M. N. R. Ashfold, R. N. Dixon, N. Little, R. J. Stickland, and C. M. Western. The $\tilde{B}^1\text{E}''$ state of ammonia: SubDoppler spectroscopy at vacuum ultraviolet energies. The Journal of Chemical Physics, 89(4):1754–1761, 1988.
- [166] M.N.R. Ashfold, S.R. Langford, R.A. Morgan, A.J. Orr-Ewing, C.M. Western, C.R. Scheper, and C.A. de Lange. Resonance enhanced multiphoton ionization (REMPI) and REMPI-photoelectron spectroscopy of ammonia. Eur. Phys. J D, 4(2):189–197, Nov 1998.
- [167] George C. Nieman and Steven D. Colson. Characterization of the \tilde{C}^2 state of ammonia observed by threephoton, gasphase spectroscopy. The Journal of Chemical Physics, 71(2):571–577, 1979.
- [168] Brian Sawyer. Cold polar molecules for novel collision experiments at low energies. PhD thesis, University of Colorado Boulder, 2010.
- [169] Franz Josef Comes. Recycling in the earth’s atmosphere: The OH radicalits importance for the chemistry of the atmosphere and the determination of its concentration. Angewandte Chemie International Edition in English, 33(18):1816–1826, 1994.
- [170] Sasho Gligorovski, Rafal Strekowski, Stephane Barbati, and Davide Vione. Environmental implications of hydroxyl radicals (OH). Chemical Reviews, 115(24):13051–13092, 2015.
- [171] John L. Bohn and Goulven Qumner. Dipolar radicals in crossed electric and magnetic fields. Molecular Physics, 111(12-13):1931–1938, 2013.
- [172] Yomay Shyur, Jason A. Bossert, and H. J. Lewandowski. Pulsed operation of a ring Stark decelerator. *Under review*.
- [173] B. Friedrich. A quasi-analytic model of a linear Stark accelerator/decelerator for polar molecules. Eur. Phys. J. D, 31(2):313–336, Nov 2004.
- [174] Yomay Shyur, N. J. Fitch, Jason A. Bossert, Terry Brown, and H. J. Lewandowski. A high-voltage amplifier for traveling-wave Stark deceleration. *Under review*.
- [175] RV Latham, editor. High Voltage Vacuum Insulation. Academic Press, London, 1995.
- [176] P. Horowitz and W. Hill. The Art of Electronics. Cambridge University Press, New York, NY, USA, 1989.

- [177] Alexander Kuhlicke, Klaus Palis, and Oliver Benson. Broadband linear high-voltage amplifier for radio frequency ion traps. *Review of Scientific Instruments*, 85(11):114707, 2014.
- [178] D. J. Larson, J. C. Bergquist, J. J. Bollinger, Wayne M. Itano, and D. J. Wineland. Sympathetic cooling of trapped ions: A laser-cooled two-species nonneutral ion plasma. *Phys. Rev. Lett.*, 57:70–73, Jul 1986.
- [179] C. J. Myatt, E. A. Burt, R. W. Ghrist, E. A. Cornell, and C. E. Wieman. Production of two overlapping bose-einstein condensates by sympathetic cooling. *Phys. Rev. Lett.*, 78:586–589, Jan 1997.
- [180] John M. Doyle, Bretislav Friedrich, Jinha Kim, and David Patterson. Buffer-gas loading of atoms and molecules into a magnetic trap. *Phys. Rev. A*, 52:R2515–R2518, Oct 1995.
- [181] Manuel Lara, John L. Bohn, Daniel Potter, Pavel Soldán, and Jeremy M. Hutson. Ultracold Rb-OH collisions and prospects for sympathetic cooling. *Phys. Rev. Lett.*, 97:183201, Nov 2006.
- [182] Manuel Lara, John L. Bohn, Daniel E. Potter, Pavel Soldán, and Jeremy M. Hutson. Cold collisions between OH and Rb: The field-free case. *Phys. Rev. A*, 75:012704, Jan 2007.
- [183] Jongseok Lim, Matthew D. Frye, Jeremy M. Hutson, and M. R. Tarbutt. Modeling sympathetic cooling of molecules by ultracold atoms. *Phys. Rev. A*, 92:053419, Nov 2015.
- [184] Christopher Ticknor and John L. Bohn. Influence of magnetic fields on cold collisions of polar molecules. *Phys. Rev. A*, 71:022709, Feb 2005.
- [185] T. V. Tscherbul, Z. Pavlovic, H. R. Sadeghpour, R. Côté, and A. Dalgarno. Collisions of trapped molecules with slow beams. *Phys. Rev. A*, 82:022704, Aug 2010.
- [186] Benjamin L. Lev, Edmund R. Meyer, Eric R. Hudson, Brian C. Sawyer, John L. Bohn, and Jun Ye. OH hyperfine ground state: From precision measurement to molecular qubits. *Phys. Rev. A*, 74:061402, Dec 2006.
- [187] Brian C. Sawyer, Benjamin L. Lev, Eric R. Hudson, Benjamin K. Stuhl, Manuel Lara, John L. Bohn, and Jun Ye. Magneto-electrostatic trapping of ground state OH molecules. *Phys. Rev. Lett.*, 98:253002, Jun 2007.
- [188] David Reens, Hao Wu, Tim Langen, and Jun Ye. Controlling spin flips of molecules in an electromagnetic trap. *Phys. Rev. A*, 96:063420, Dec 2017.
- [189] Moritz Kirste, Boris G. Sartakov, Melanie Schnell, and Gerard Meijer. Nonadiabatic transitions in electrostatically trapped ammonia molecules. *Phys. Rev. A*, 79:051401, May 2009.
- [190] T. E. Wall, S. K. Tokunaga, E. A. Hinds, and M. R. Tarbutt. Nonadiabatic transitions in a Stark decelerator. *Phys. Rev. A*, 81:033414, Mar 2010.
- [191] L. Paul Parazzoli. Electric-field affected low-energy collisions between co-trapped ammonia and rubidium. PhD thesis, University of Colorado Boulder.
- [192] A.J.C. Varandas and A.I. Voronin. Calculation of the asymptotic interaction and modelling of the potential energy curves of OH and OH⁺. *Chemical Physics*, 194(1):91 – 100, 1995.

- [193] Esther de Beer, M. P. Koopmans, C. A. de Lange, Yumin Wang, and W. A. Chupka. (2+1) resonance-enhanced multiphoton ionization photoelectron spectroscopy of the OH radical. The Journal of Chemical Physics, 94(12):7634–7639, 1991.
- [194] Margaret E. Greenslade, Marsha I. Lester, Dragana . Radenovi, Andr J. A. van Roij, and David H. Parker. (2+1) resonance-enhanced ionization spectroscopy of a state-selected beam of OH radicals. The Journal of Chemical Physics, 123(7):074309, 2005.
- [195] R. Forster, H. Hippler, K. Hoyer mann, G. Rohde, and Lawrence B. Harding. REMPI mass spectrum of the OH radical in the gas phase. Chemical Physics Letters, 183(6):465 – 470, 1991.
- [196] Joseph M. Beames, Fang Liu, Marsha I. Lester, and Craig Murray. Communication: A new spectroscopic window on hydroxyl radicals using UV + VUV resonant ionization. The Journal of Chemical Physics, 134(24):241102, 2011.
- [197] Amy M. Green, Fang Liu, and Marsha I. Lester. UV + V UV double-resonance studies of autoionizing Rydberg states of the hydroxyl radical. The Journal of Chemical Physics, 144(18):184311, 2016.
- [198] Manuel Lara, Benjamin L. Lev, and John L. Bohn. Loss of molecules in magneto-electrostatic traps due to nonadiabatic transitions. Phys. Rev. A, 78:033433, Sep 2008.
- [199] M. Bhattacharya, Z. Howard, and M. Kleinert. Ground-state OH molecule in combined electric and magnetic fields: Analytic solution of the effective hamiltonian. Phys. Rev. A, 88:012503, Jul 2013.
- [200] Nicholas P. Lockyer and John C. Vickerman. Single photon ionisation mass spectrometry using laser-generated vacuum ultraviolet photons. Laser Chem, 17:139–159, 1997.
- [201] Y. J. Shi, S. Consta, A. K. Das, B. Mallik, D. Lacey, and R. H. Lipson. A 118 nm vacuum ultraviolet laser/time-of-flight mass spectroscopic study of methanol and ethanol clusters in the vapor phase. The Journal of Chemical Physics, 116(16), 2002.
- [202] R Miles and S Harris. Optical third-harmonic generation in alkali metal vapors. Quantum Electronics, IEEE Journal of, 9(4):470–484, 1973.
- [203] Gary C Bjorklund. Effects of focusing on third-order nonlinear processes in isotropic media. Quantum Electronics, IEEE Journal of, 11(6):287–296, 1975.
- [204] JF Ward and GHC New. Optical third harmonic generation in gases by a focused laser beam. Physical Review, 185(1):57, 1969.
- [205] R. Mahon, T. McIlrath, V. Myerscough, and D. Koopman. Third-harmonic generation in argon, krypton, and xenon: Bandwidth limitations in the vicinity of lyman- α ;. IEEE Journal of Quantum Electronics, 15(6):444–451, Jun 1979.
- [206] A. Kramida, Yu. Ralchenko, J. Reader, and and NIST ASD Team. NIST Atomic Spectra Database (ver. 5.3), [Online]. Available: <http://physics.nist.gov/asd> [2016, January 29]. National Institute of Standards and Technology, Gaithersburg, MD., 2015.

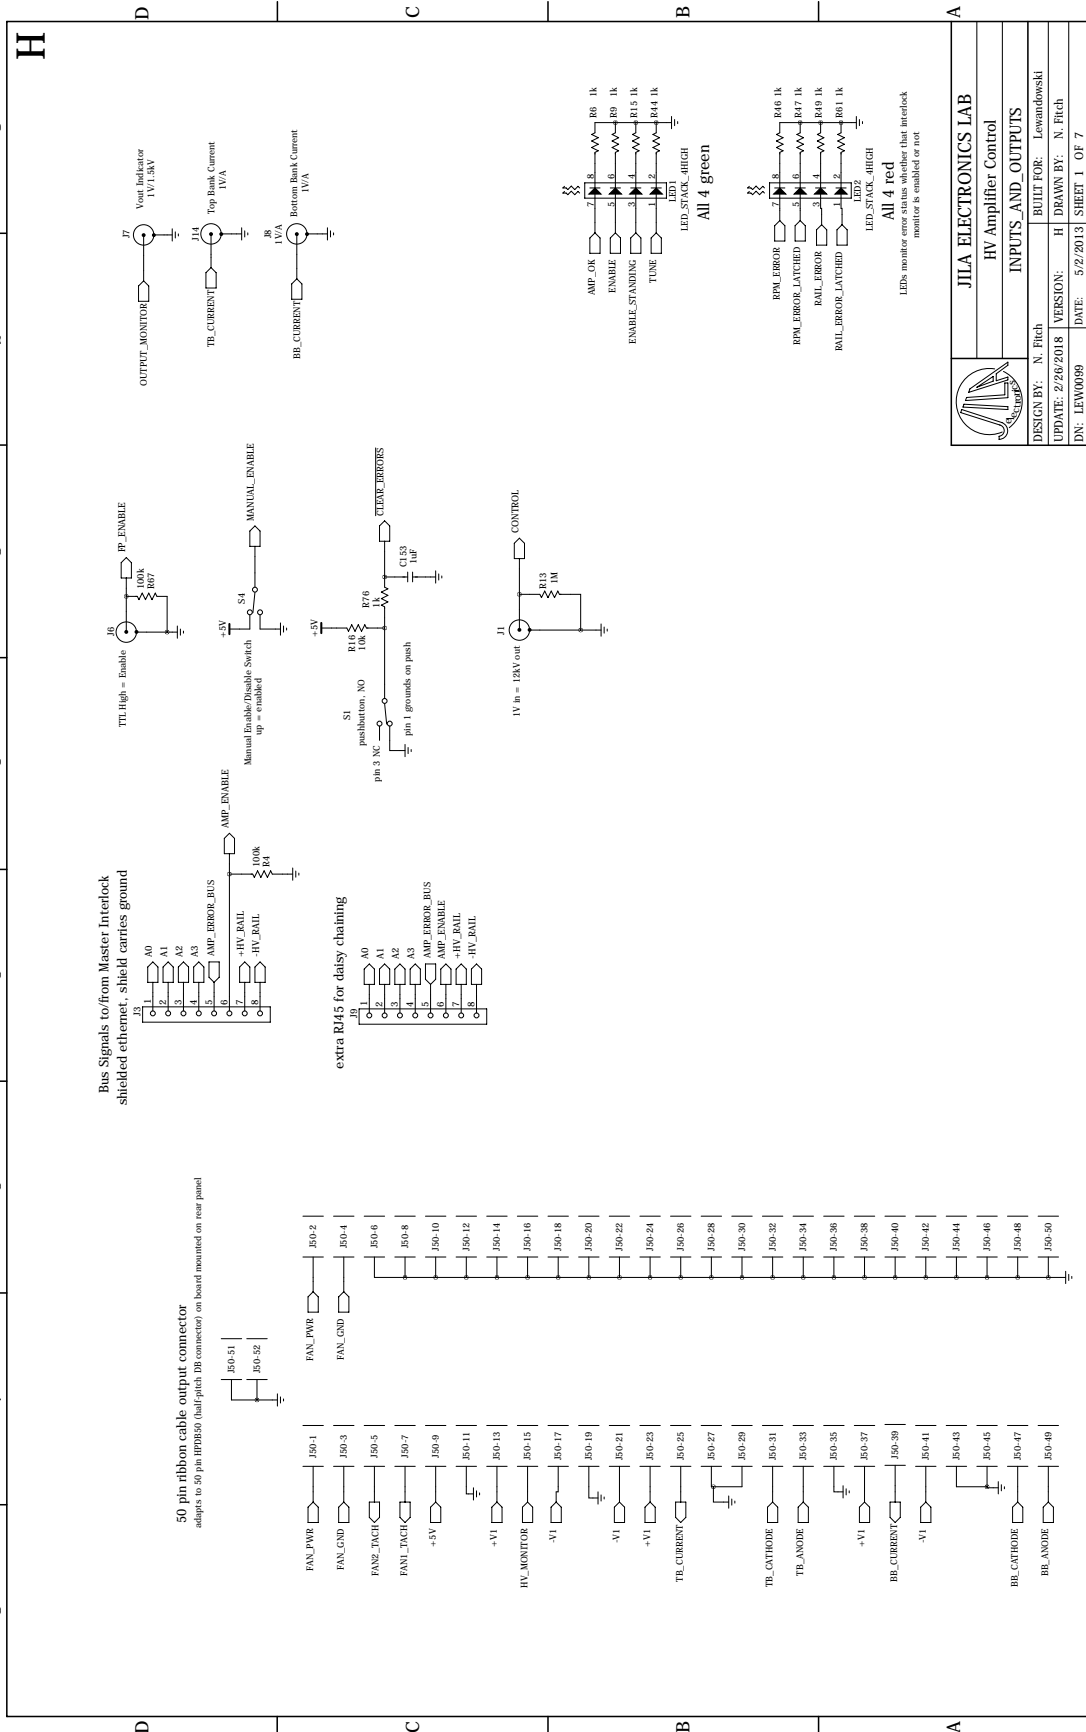
- [207] R. S. Finn and J. F. Ward. dc-induced optical second-harmonic generation in the inert gases. Phys. Rev. Lett., 26:285–289, Feb 1971.
- [208] A. Flusberg, T. Mossberg, and S. R. Hartmann. Optical difference-frequency generation in atomic thallium vapor. Phys. Rev. Lett., 38:59–62, Jan 1977.
- [209] M. Matsuoka, H. Nakatsuka, H. Uchiki, and M. Mitsunaga. Optical second-harmonic generation in gases: "rotation" of quadrupole moment in magnetic field. Phys. Rev. Lett., 38:894–898, Apr 1977.
- [210] T. Mossberg, A. Flusberg, and S.R. Hartmann. Optical second-harmonic generation in atomic thallium vapor. Optics Communications, 25(1):121 – 124, 1978.
- [211] Kenzo Miyazaki, Takuzo Sato, and Hiroshi Kashiwagi. Spontaneous-field-induced optical second-harmonic generation in atomic vapors. Phys. Rev. Lett., 43:1154–1157, Oct 1979.
- [212] W. Jamroz, P. E. LaRocque, and B. P. Stoicheff. Resonantly enhanced second-harmonic generation in zinc vapor. Opt. Lett., 7(4):148–150, Apr 1982.
- [213] Michelle S. Malcuit, Robert W. Boyd, William V. Davis, and Kazimierz Rzaewski. Anomalies in optical harmonic generation using high-intensity laser radiation. Phys. Rev. A, 41:3822–3825, Apr 1990.
- [214] Y. J. Shi, W. Al-Basheer, and R. I. Thompson. Two-photon resonant second harmonic generation in atomic xenon. The Journal of Chemical Physics, 130(9), 2009.
- [215] L. Marmet, K. Hakuta, and B. P. Stoicheff. Second-harmonic generation in atomic hydrogen induced by a charge-separation field. J. Opt. Soc. Am. B, 9(7):1038–1046, Jul 1992.
- [216] Donald S. Bethune, Robert W. Smith, and Y. R. Shen. Sum-frequency generation via a resonant quadrupole transition in sodium. Phys. Rev. A, 17:277–292, Jan 1978.
- [217] D. S. Bethune. Optical second-harmonic generation in atomic vapors with focused beams. Phys. Rev. A, 23:3139–3151, Jun 1981.

Appendix A


Circuit Diagrams

This appendix contains the full updated circuit schematics, updated from [95] for the amplifiers discussed in Chapter 5. The circuits are included in the following order (circuit titles are found in the lower righthand corner of each circuit):


- Control board: this board creates the driving signals for both banks of one amplifier. The PID feedback loop for each amplifier is located on this board.
- Amplifier bank: this is one half of one amplifier. One bank consists of four circuit board the main board containing the follow and nine slaves, a high-voltage FET wall, a resistor-divider wall, and a protection diode wall. Two banks make up one amplifier.
- Amplifier breakout: this board interfaces between the amplifier (via ethernet cables) and the control board (via a HPDB50 cable).
- Master interlock: this board conducts global error monitoring including cooling water, power supplies, and amplifiers (through the error bus)
- Master interlock breakout: this board converts between two DB25 connectors for the Spellman high-voltage supplies and two ethernet error bus connectors to one 50 pin ribbon cable that connects to the master interlock




JILA ELECTRONICS LAB	
HV Amplifier Control	
INPUTS AND OUTPUTS	
DESIGN BY: N. Fitch	BUILT FOR: Lewandowski
UPDATE: 2/26/2018	VERSION: H1
DRAWN BY: N. Fitch	DATE: 5/2/2013
IDN: LEW0069	SHEET 1 OF 7

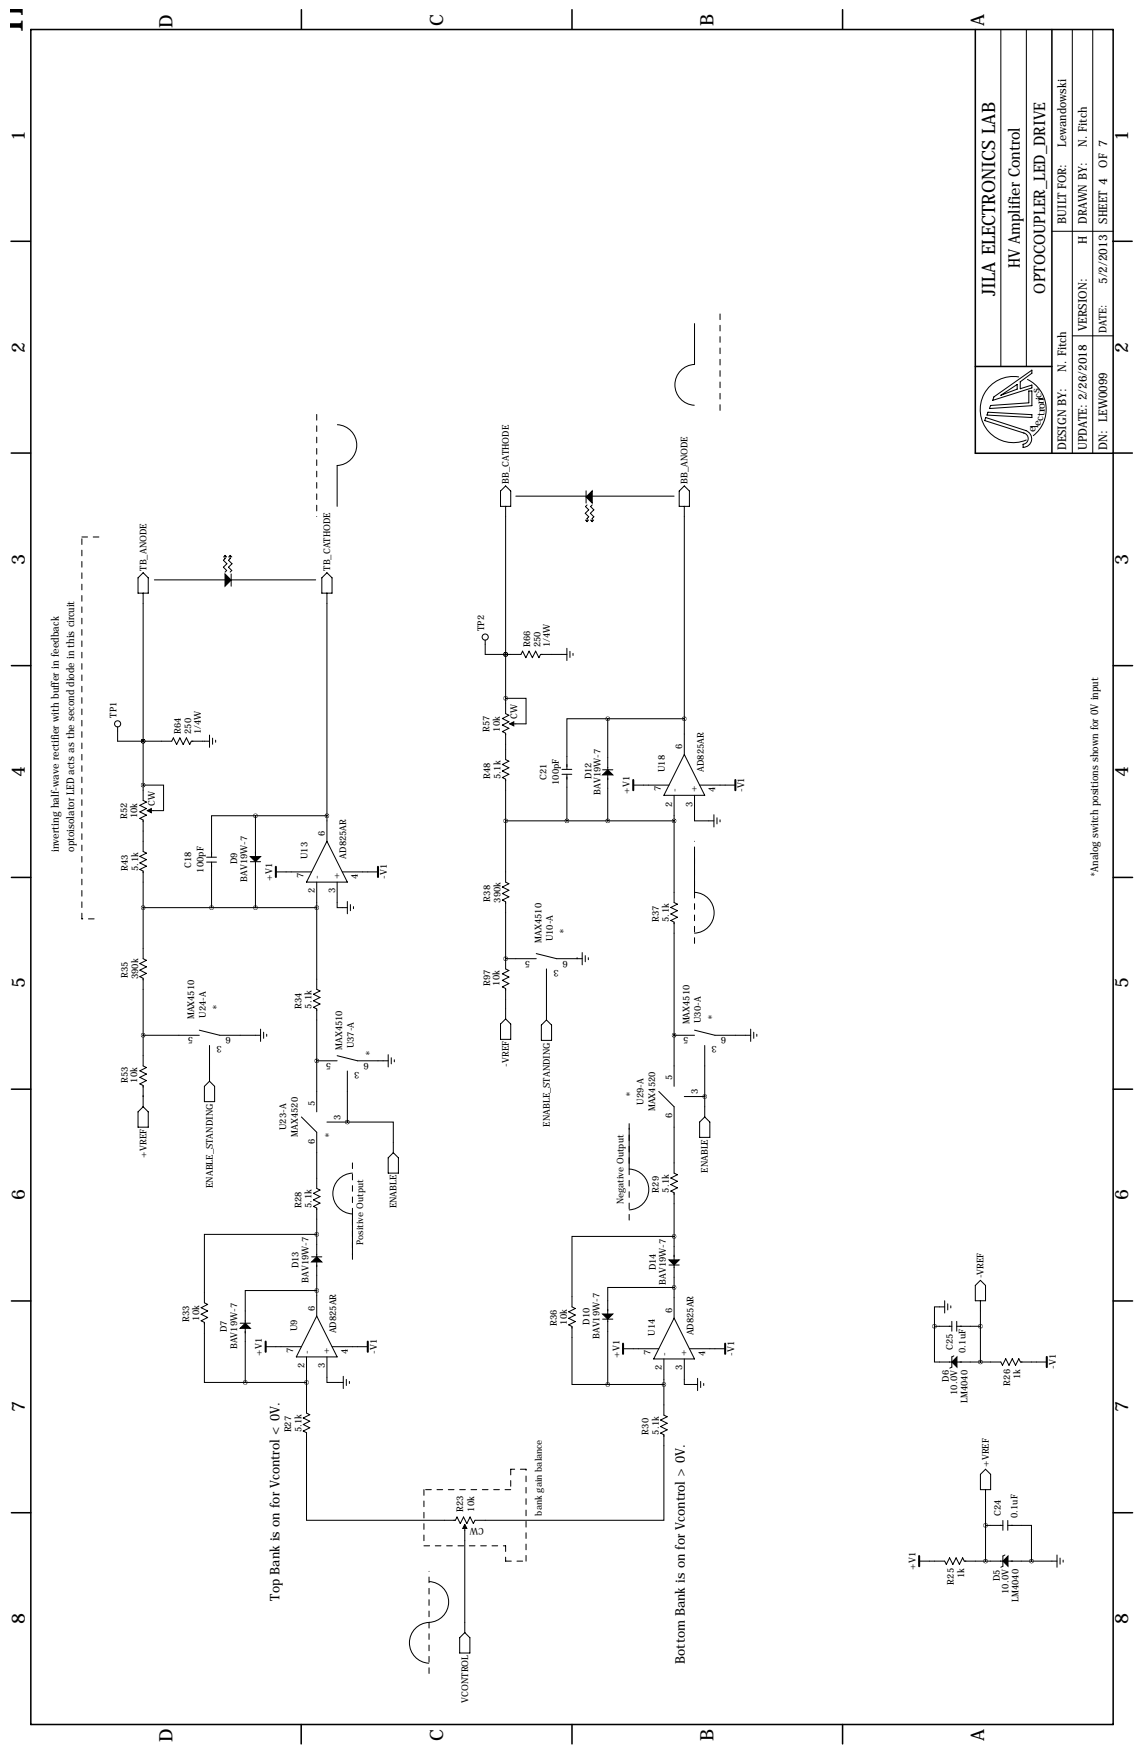



TTL enable signal can come from the master interlock board (rear panel ethernet) or from front panel smu (for testing)

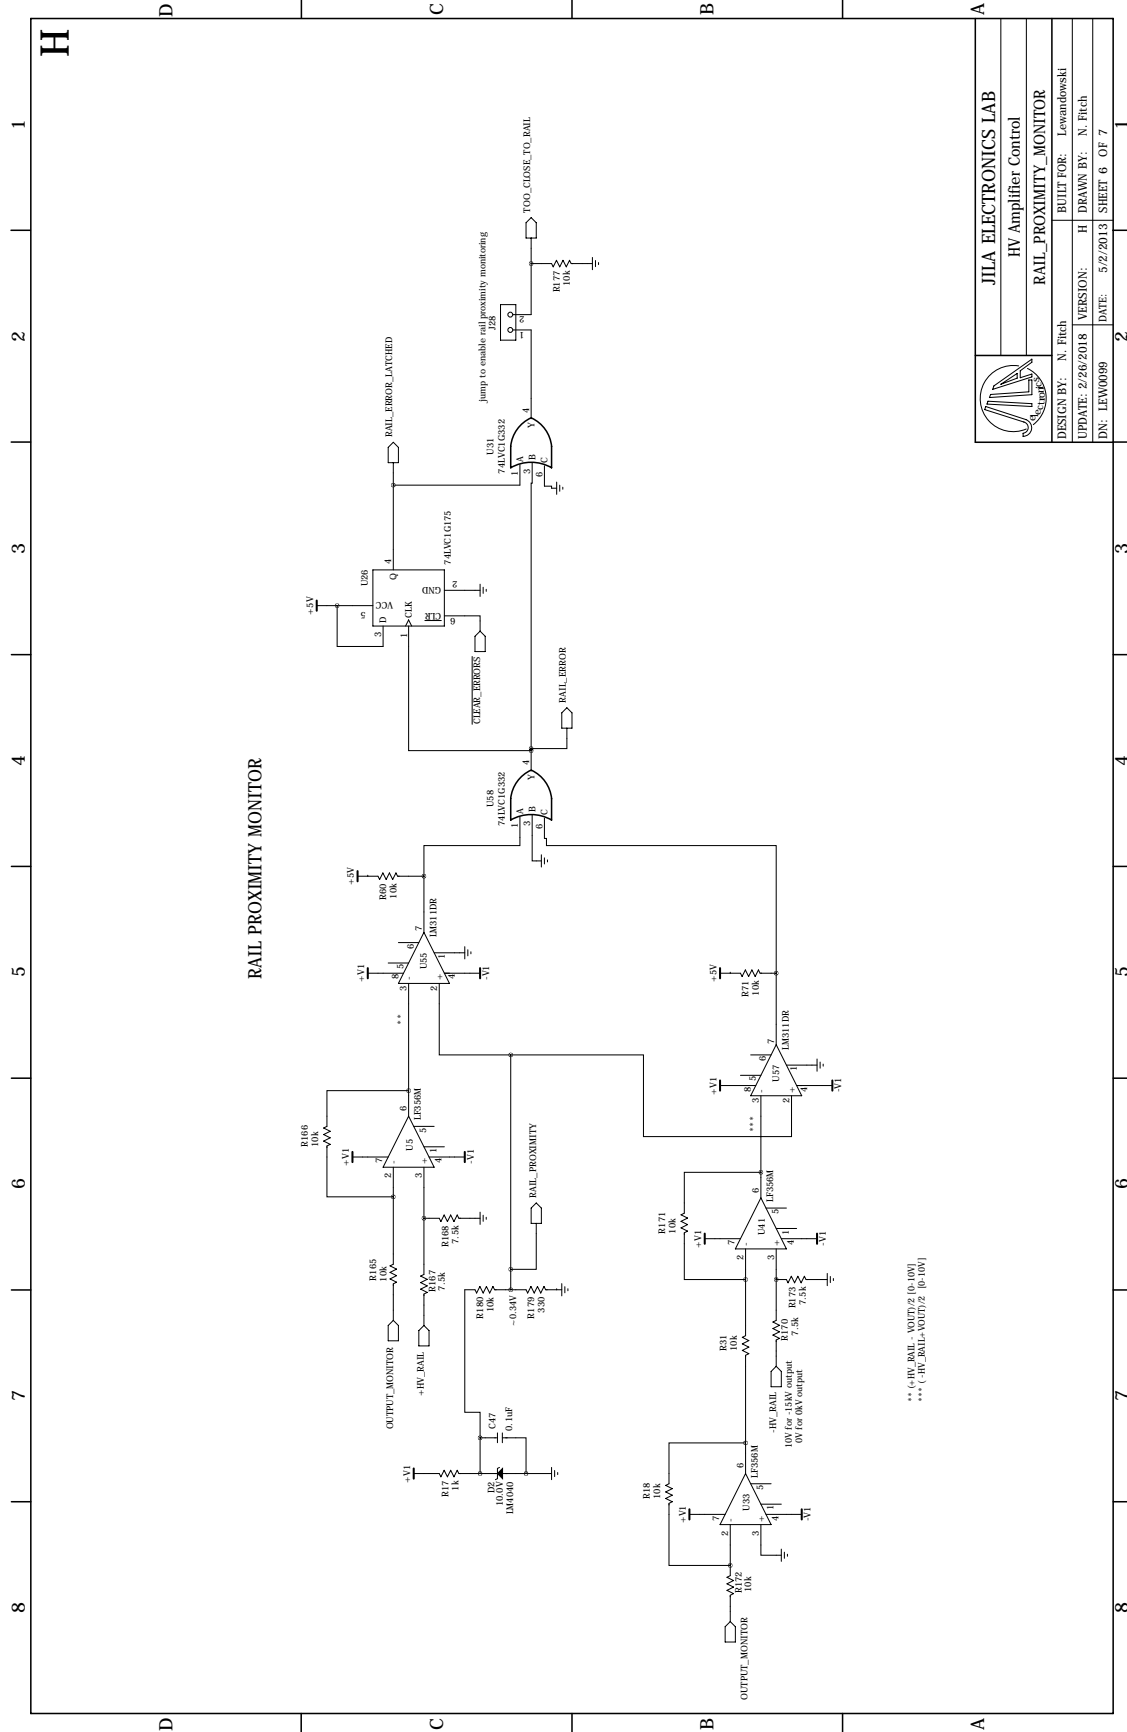
	
JILA ELECTRONICS LAB HV Amplifier Control	
LOGIC	
DESIGN BY: N. Fitch	BUILT FOR: Lewandowski
UPDATE: 2/26/2018	VERSION: H
DATE: 5/2/2013	DRAWN BY: N. Fitch
IDN: LEW0009	SHEET 2 OF 7




JILA ELECTRONICS LAB	
HV Amplifier Control	
FEEDBACK_LOOP	
DESIGN BY: N. Fitch	BUILT FOR: Lewandowski
UPDATE: 2/26/2018	VERSION: H
DRAWN BY: N. Fitch	DATE: 5/2/2013
DN: LEW0069	SHEET 3 OF 7




		JILA ELECTRONICS LAB HV Amplifier Control	
DESIGN BY: N. Fitch UPDATE: 2/26/2018 DN: LEW0009		BUILT FOR: Lewandowski DRAWN BY: N. Fitch DATE: 5/2/2013 SHEET 4 OF 7	
OPTOCOUPLER_LED_DRIVE		2	1



RAIL PROXIMITY MONITOR

	
JILA ELECTRONICS LAB HV Amplifier Control	
RAIL_PROXIMITY_MONITOR	
DESIGN BY: N. Fitch	BUILT FOR: Lewandowski
UPDATE: 2/26/2018	VERSION: H
DRAWN BY: N. Fitch	DATE: 5/2/2013
IDN: LEW0069	SHEET 6 OF 7

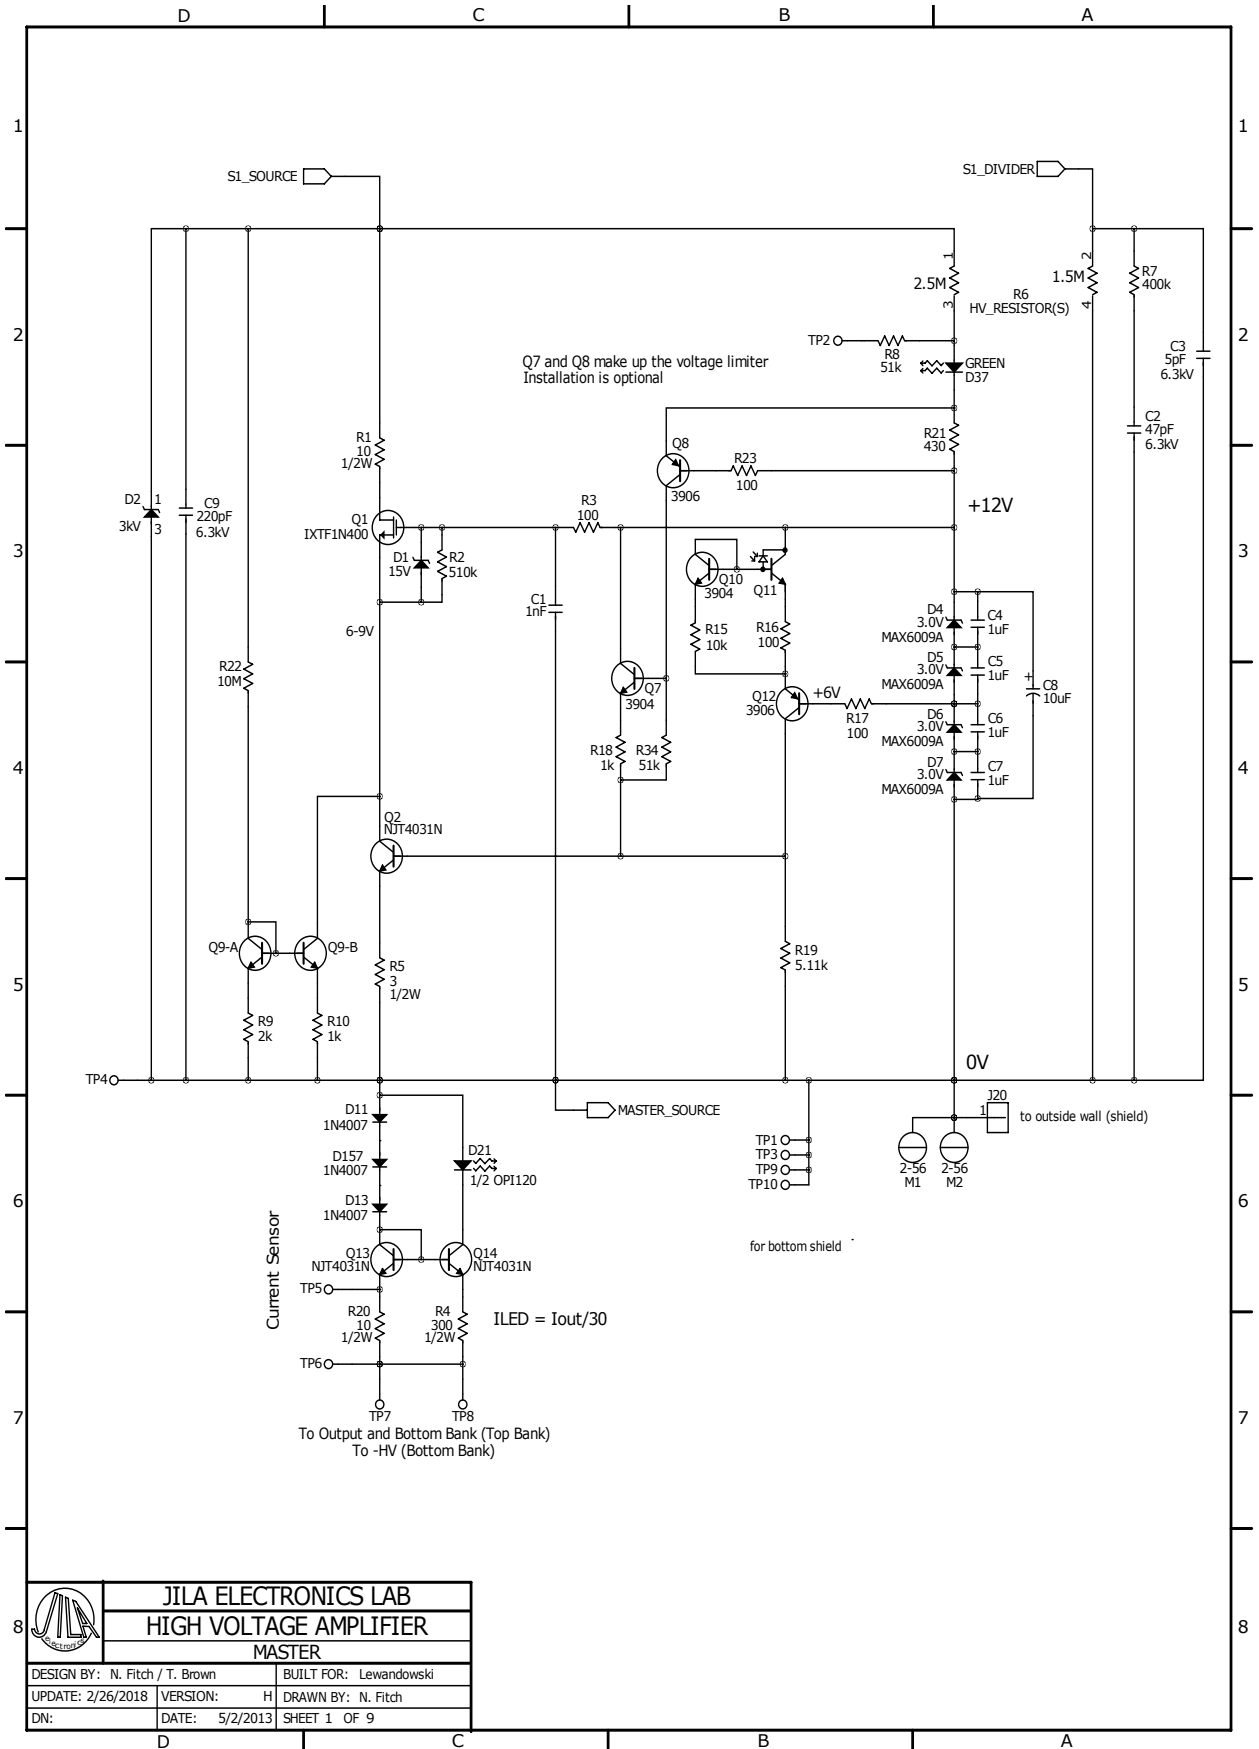
*** (HW_RAIL_MONITOR) (P.001) ***



JILA ELECTRONICS LAB	
HV Amplifier Control	
POWER CONNECTIONS	
DESIGN BY: N. Fitch	BUILT FOR: Lewandowski
UPDATE: 2/26/2018	VERSION: H
DN: LEW0009	DATE: 5/2/2013
	DRAWN BY: N. Fitch
	SHEET 7 OF 7

H
1
2
3
4
5
6
7
8

A
B
C
D
I

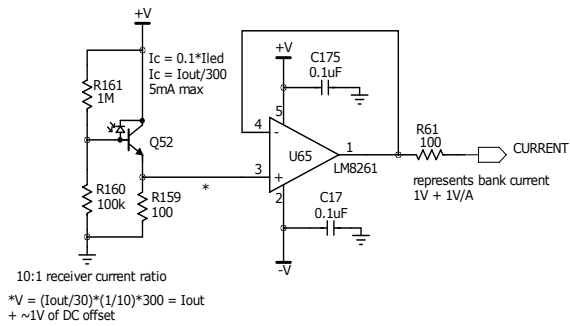




JILA ELECTRONICS LAB
HIGH VOLTAGE AMPLIFIER

MASTER

DESIGN BY: N. Fitch / T. Brown	BUILT FOR: Lewandowski
UPDATE: 2/26/2018	VERSION: H
DATE: 5/2/2013	DRAWN BY: N. Fitch
DN:	SHEET 1 OF 9

OPTOCOPLER CONNECTIONS (elevated off board)



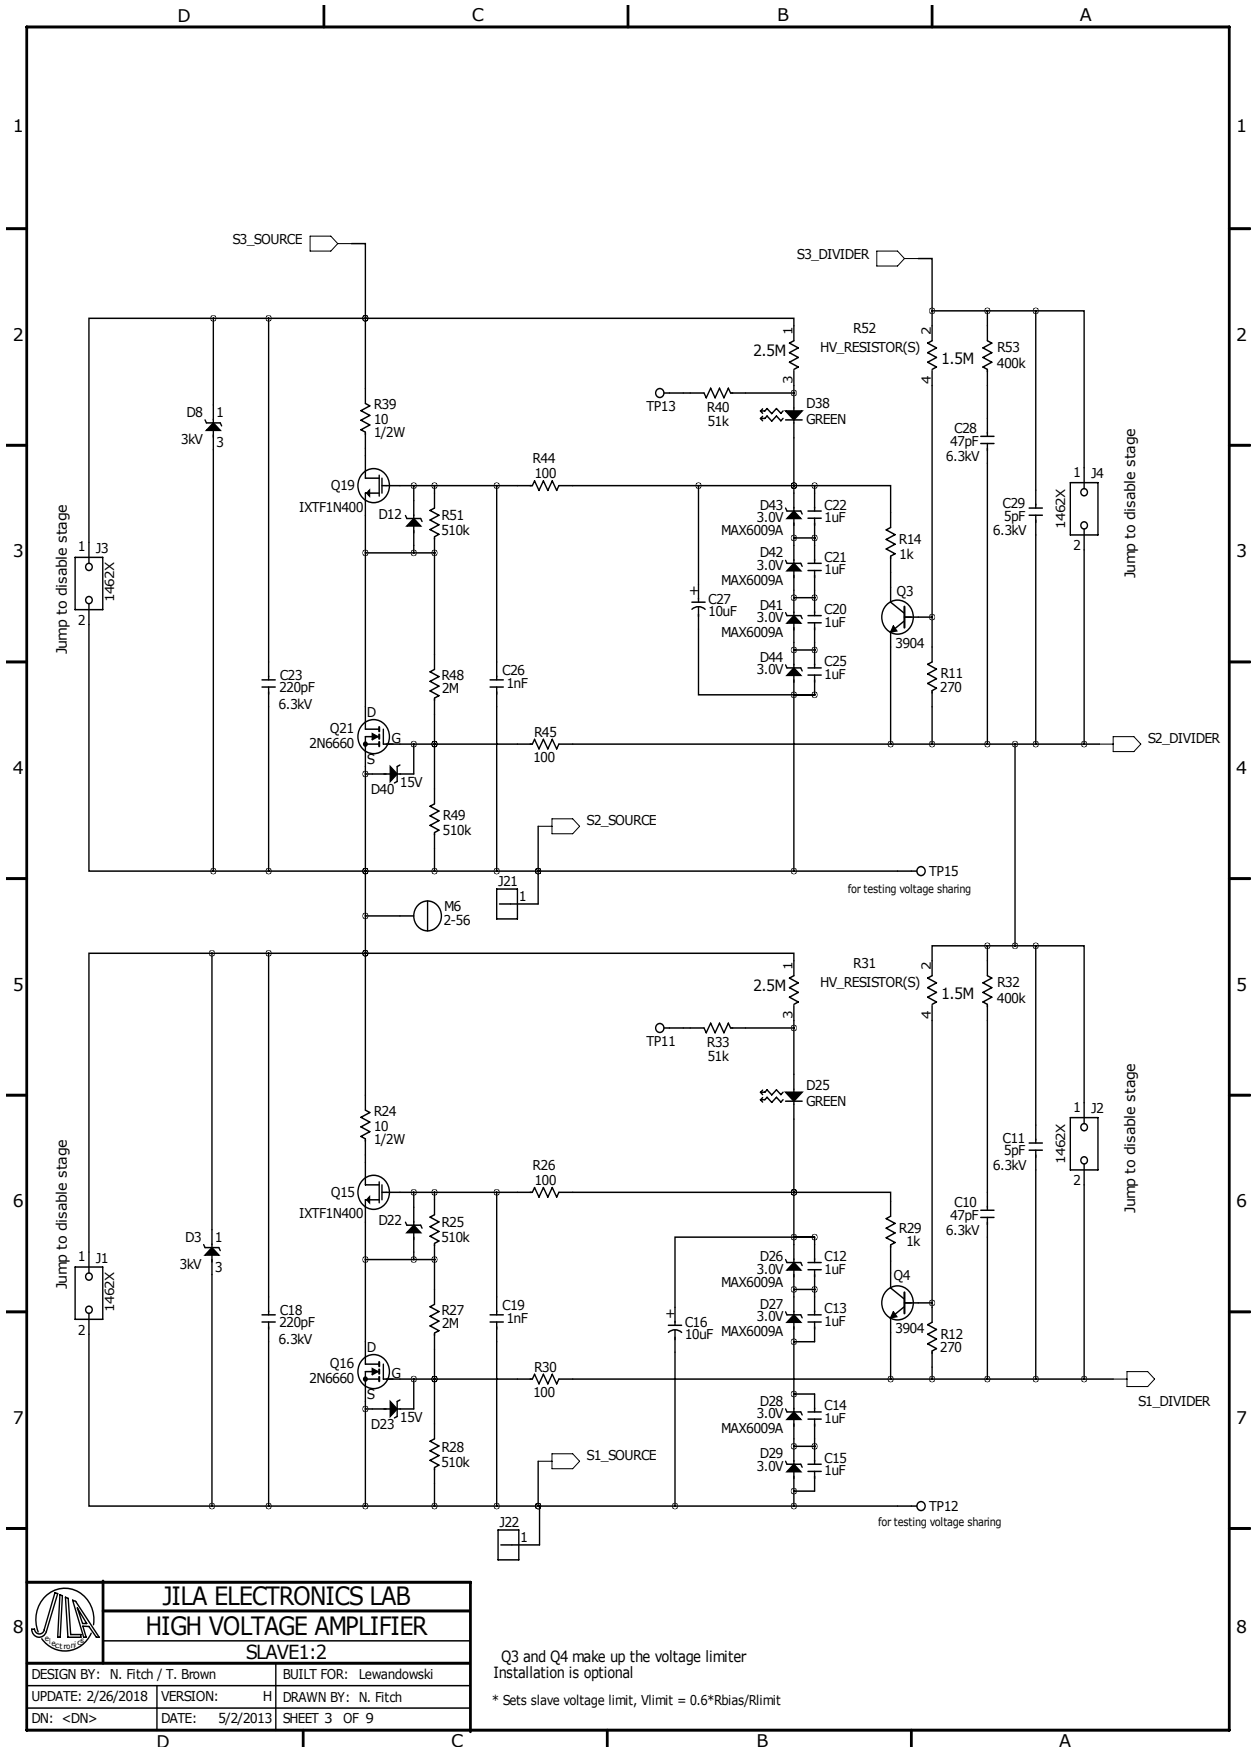
OPI120 phototransistor is insensitive to LED currents below ~250uA
 Since the LED drive is $I_{out}/30$, the current sensor is insensitive to I_{out} less than 7.5mA or so

OPI120 phototransistor case (shield) is connected to the collector
 OPI120 LED case (shield) is connected to the anode
 Common collector configuration chosen so that phototransistor body/shield is a low impedance point



JILA ELECTRONICS LAB
HIGH VOLTAGE AMPLIFIER
OPTOCOPLERS

DESIGN BY: N. Fitch / T. Brown	BUILT FOR: Lewandowski
UPDATE: 2/26/2018	VERSION: H DRAWN BY: N. Fitch
DN: <DN>	DATE: 5/2/2013 SHEET 2 OF 9



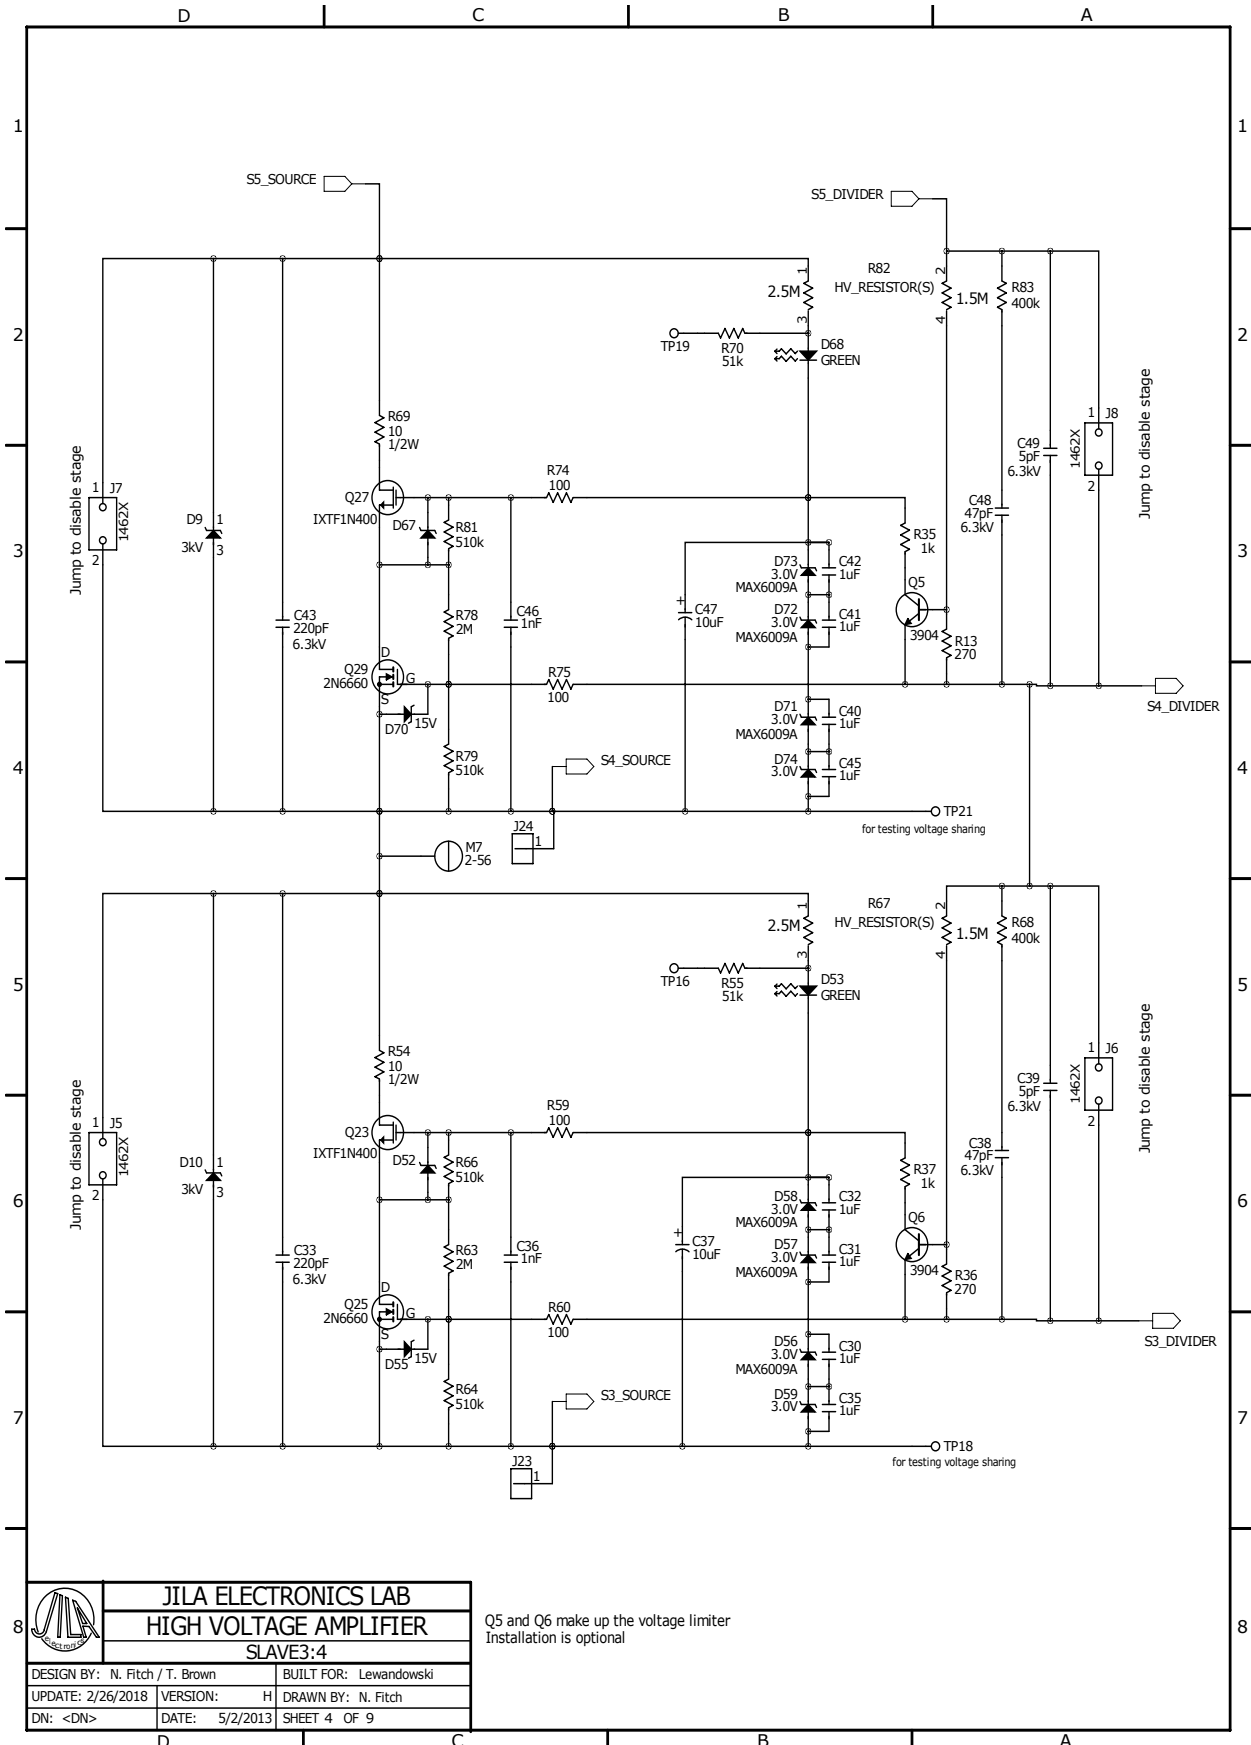
JILA ELECTRONICS LAB
HIGH VOLTAGE AMPLIFIER

SLAVE1:2

DESIGN BY: N. Fitch / T. Brown	BUILT FOR: Lewandowski
UPDATE: 2/26/2018	VERSION: H
DATE: 5/2/2013	DRAWN BY: N. Fitch
DN: <DN>	SHEET 3 OF 9

Q3 and Q4 make up the voltage limiter
 Installation is optional

* Sets slave voltage limit, $V_{limit} = 0.6 \cdot R_{bias} / R_{limit}$

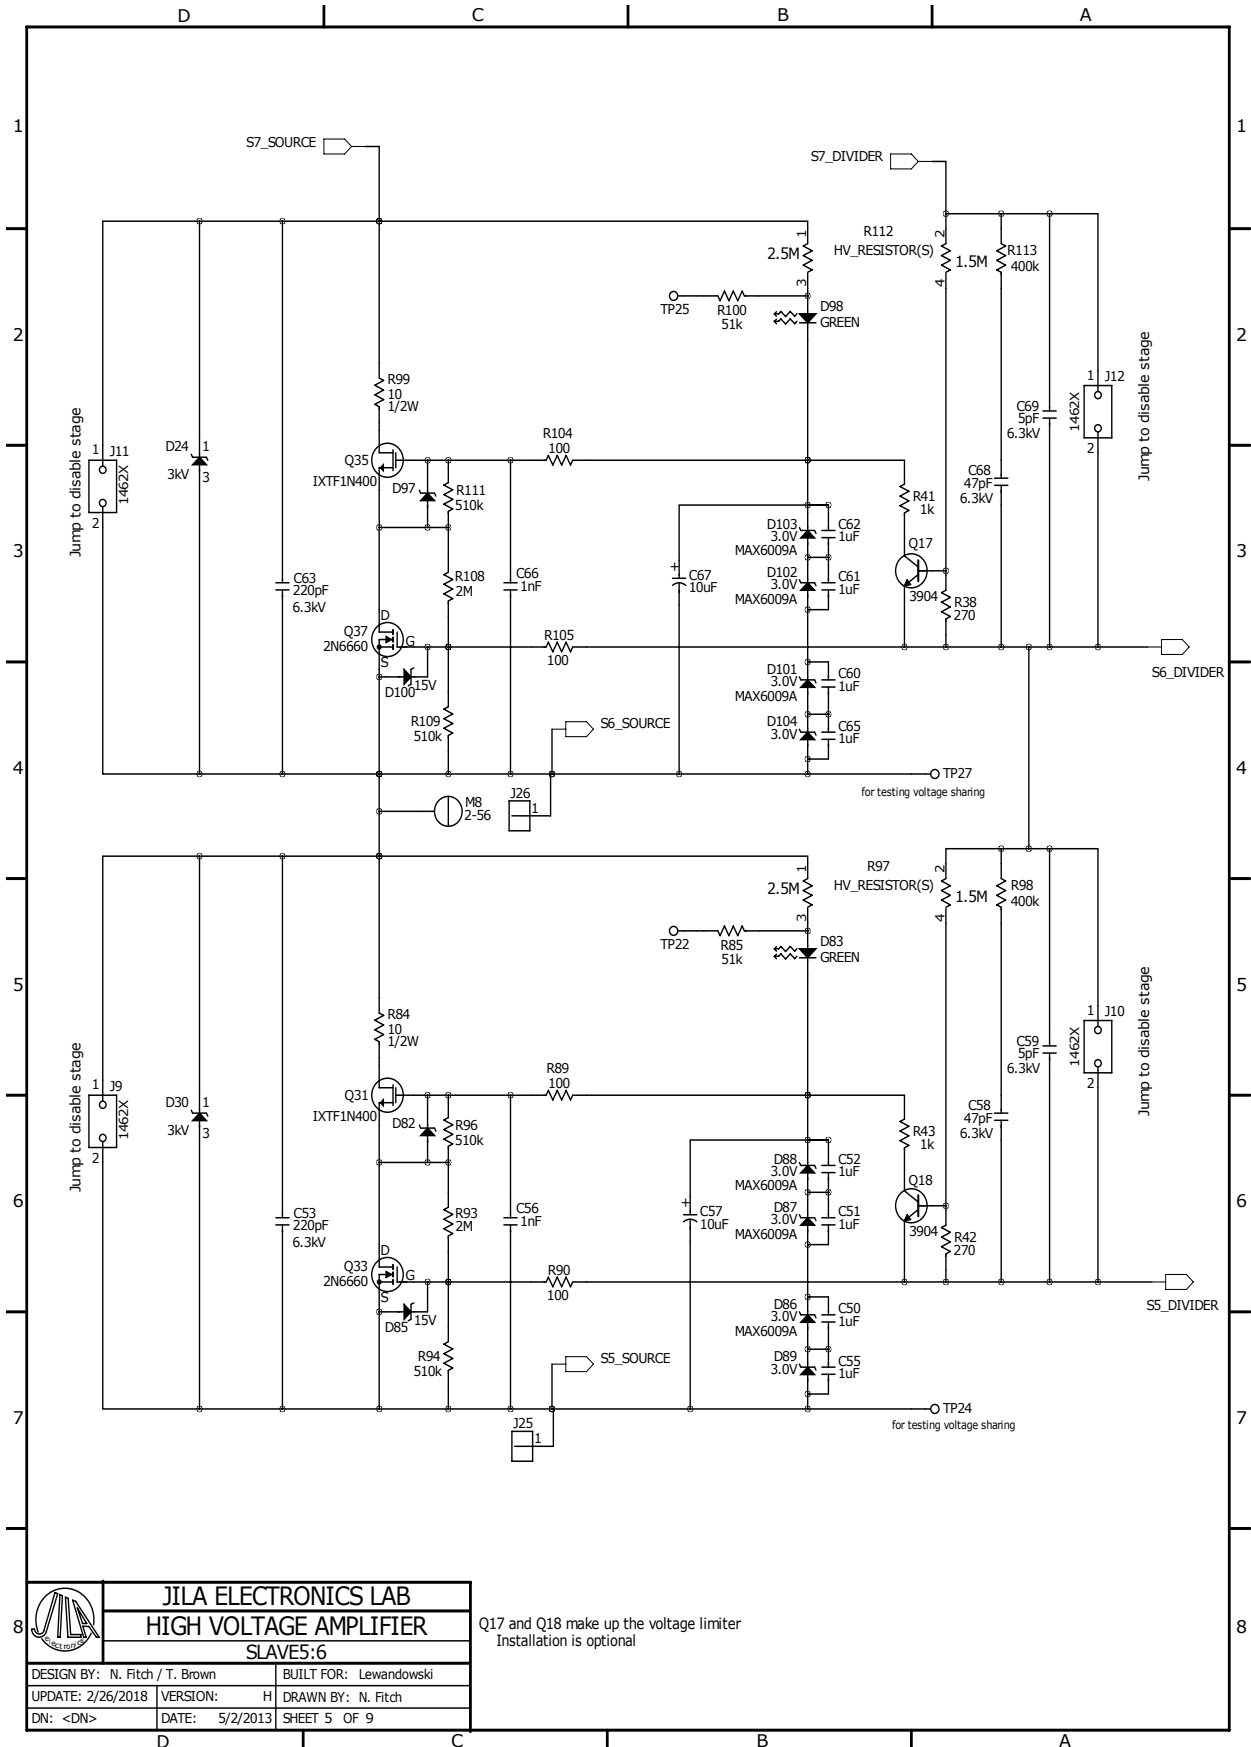


JILA ELECTRONICS LAB
HIGH VOLTAGE AMPLIFIER
SLAVE3:4

Q5 and Q6 make up the voltage limiter
 Installation is optional

DESIGN BY: N. Fitch / T. Brown	BUILT FOR: Lewandowski
UPDATE: 2/26/2018	VERSION: H
DN: <DN>	DATE: 5/2/2013
	SHEET 4 OF 9

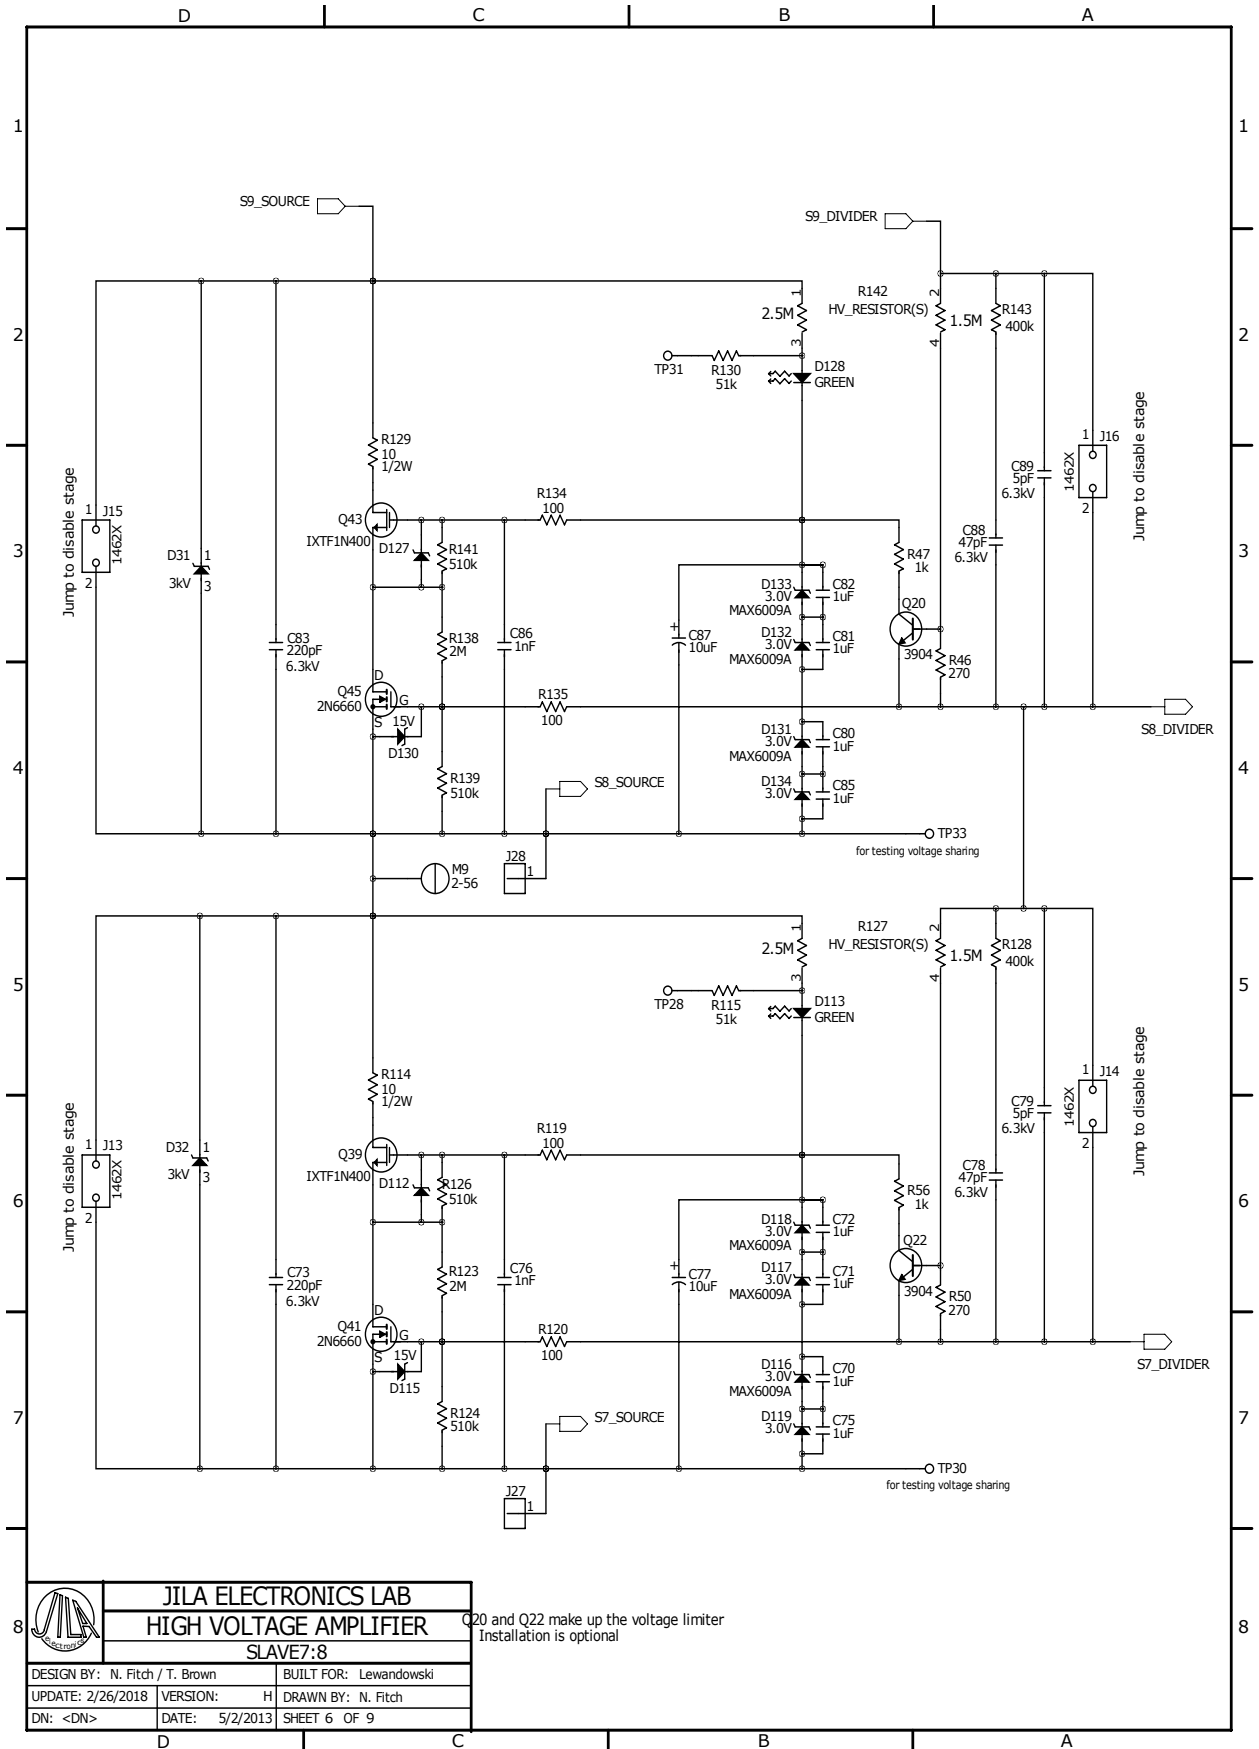
D C B A



JILA ELECTRONICS LAB
HIGH VOLTAGE AMPLIFIER
 SLAVE5:6

Q17 and Q18 make up the voltage limiter
 Installation is optional

DESIGN BY: N. Fitch / T. Brown	BUILT FOR: Lewandowski
UPDATE: 2/26/2018	VERSION: H
DATE: 5/2/2013	DRAWN BY: N. Fitch
DN: <DN>	SHEET 5 OF 9



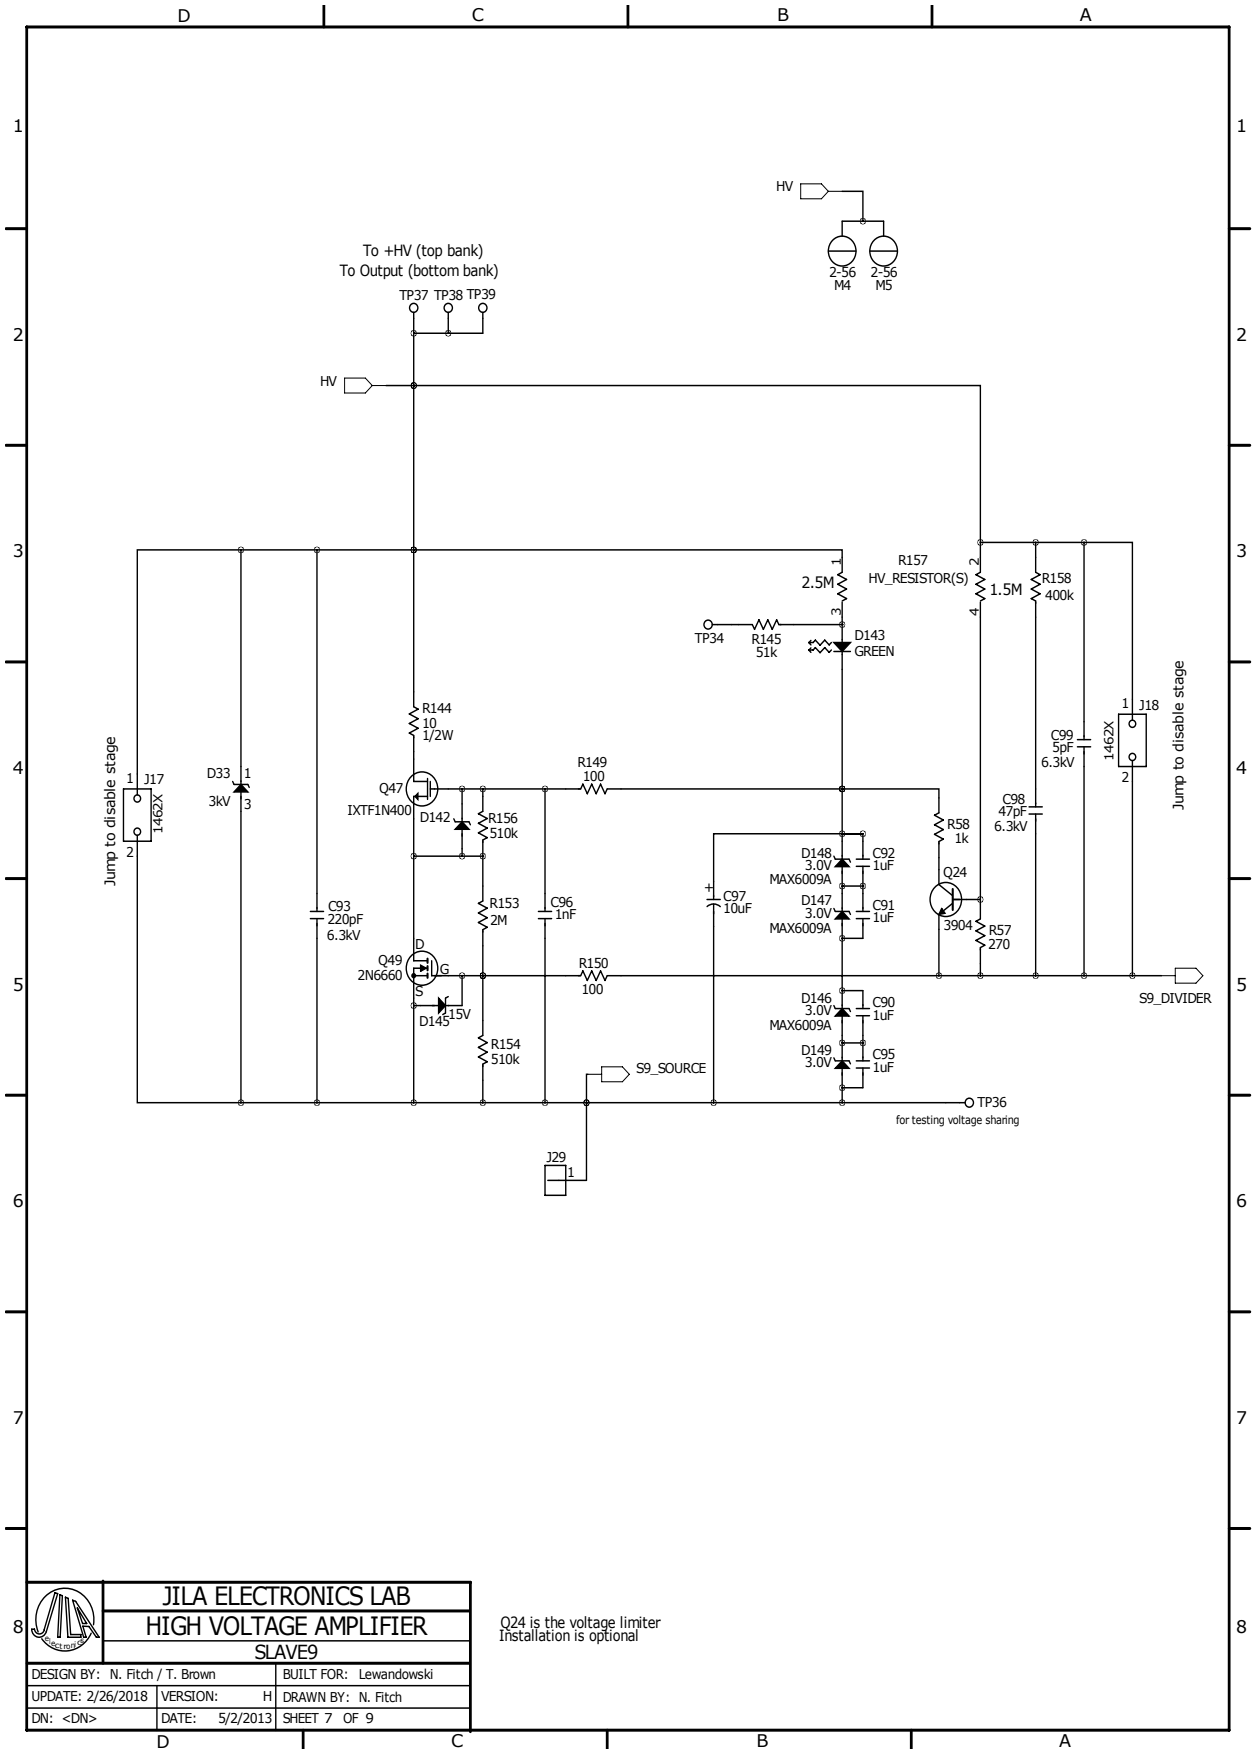
JILA ELECTRONICS LAB
HIGH VOLTAGE AMPLIFIER
SLAVE7:8


Q20 and Q22 make up the voltage limiter
 Installation is optional

DESIGN BY: N. Fitch / T. Brown	BUILT FOR: Lewandowski
UPDATE: 2/26/2018	VERSION: H
DN: <DN>	DATE: 5/2/2013
	SHEET 6 OF 9


8

8



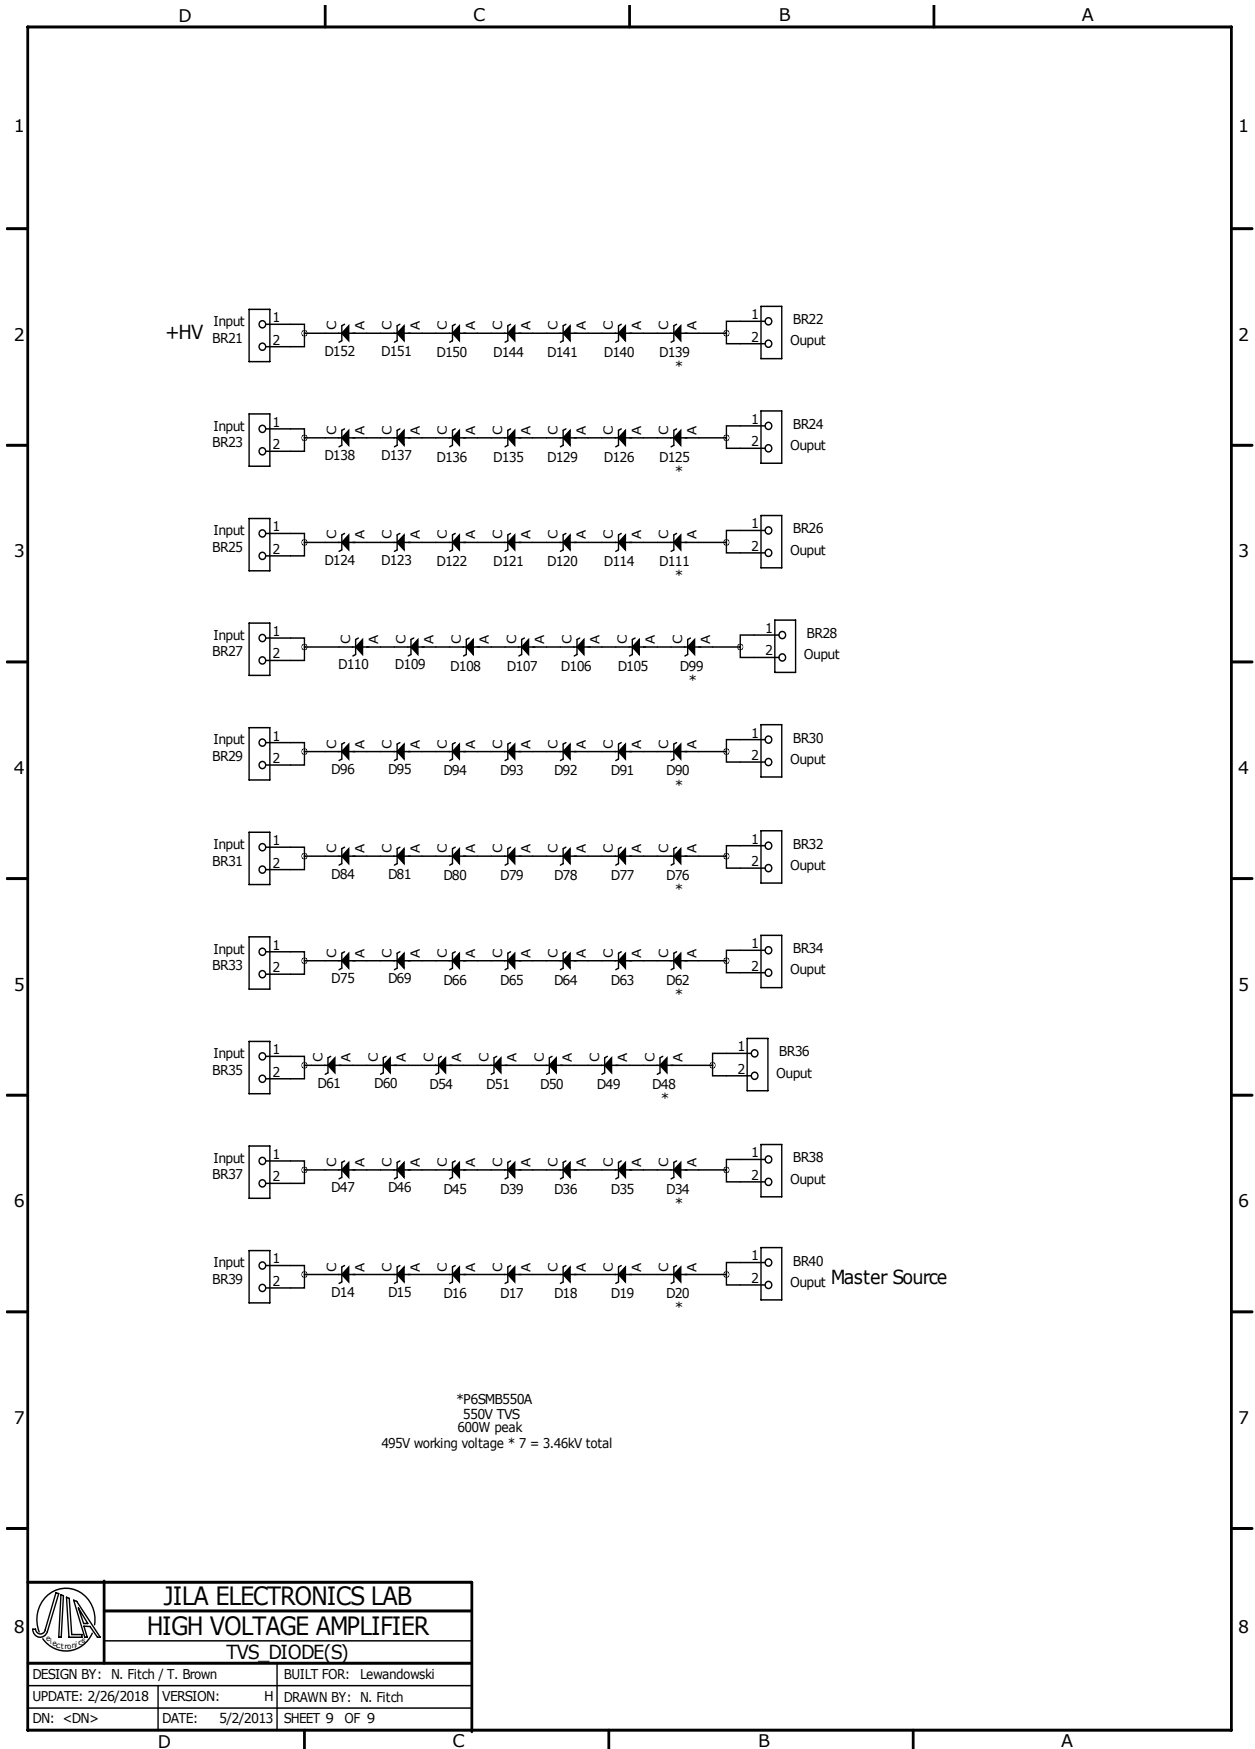
		
DESIGN BY: N. Fitch / T. Brown	BUILT FOR: Lewandowski	
UPDATE: 2/26/2018	VERSION: H	DRAWN BY: N. Fitch
DN: <DN>	DATE: 5/2/2013	SHEET 7 OF 9

Q24 is the voltage limiter
Installation is optional





JILA ELECTRONICS LAB
HIGH VOLTAGE AMPLIFIER
BIAS RESISTOR


DESIGN BY: N. Fitch / T. Brown	BUILT FOR: Lewandowski
UPDATE: 2/26/2018	VERSION: H
DATE: 5/2/2013	DRAWN BY: N. Fitch
DN: <DN>	SHEET 8 OF 9




H



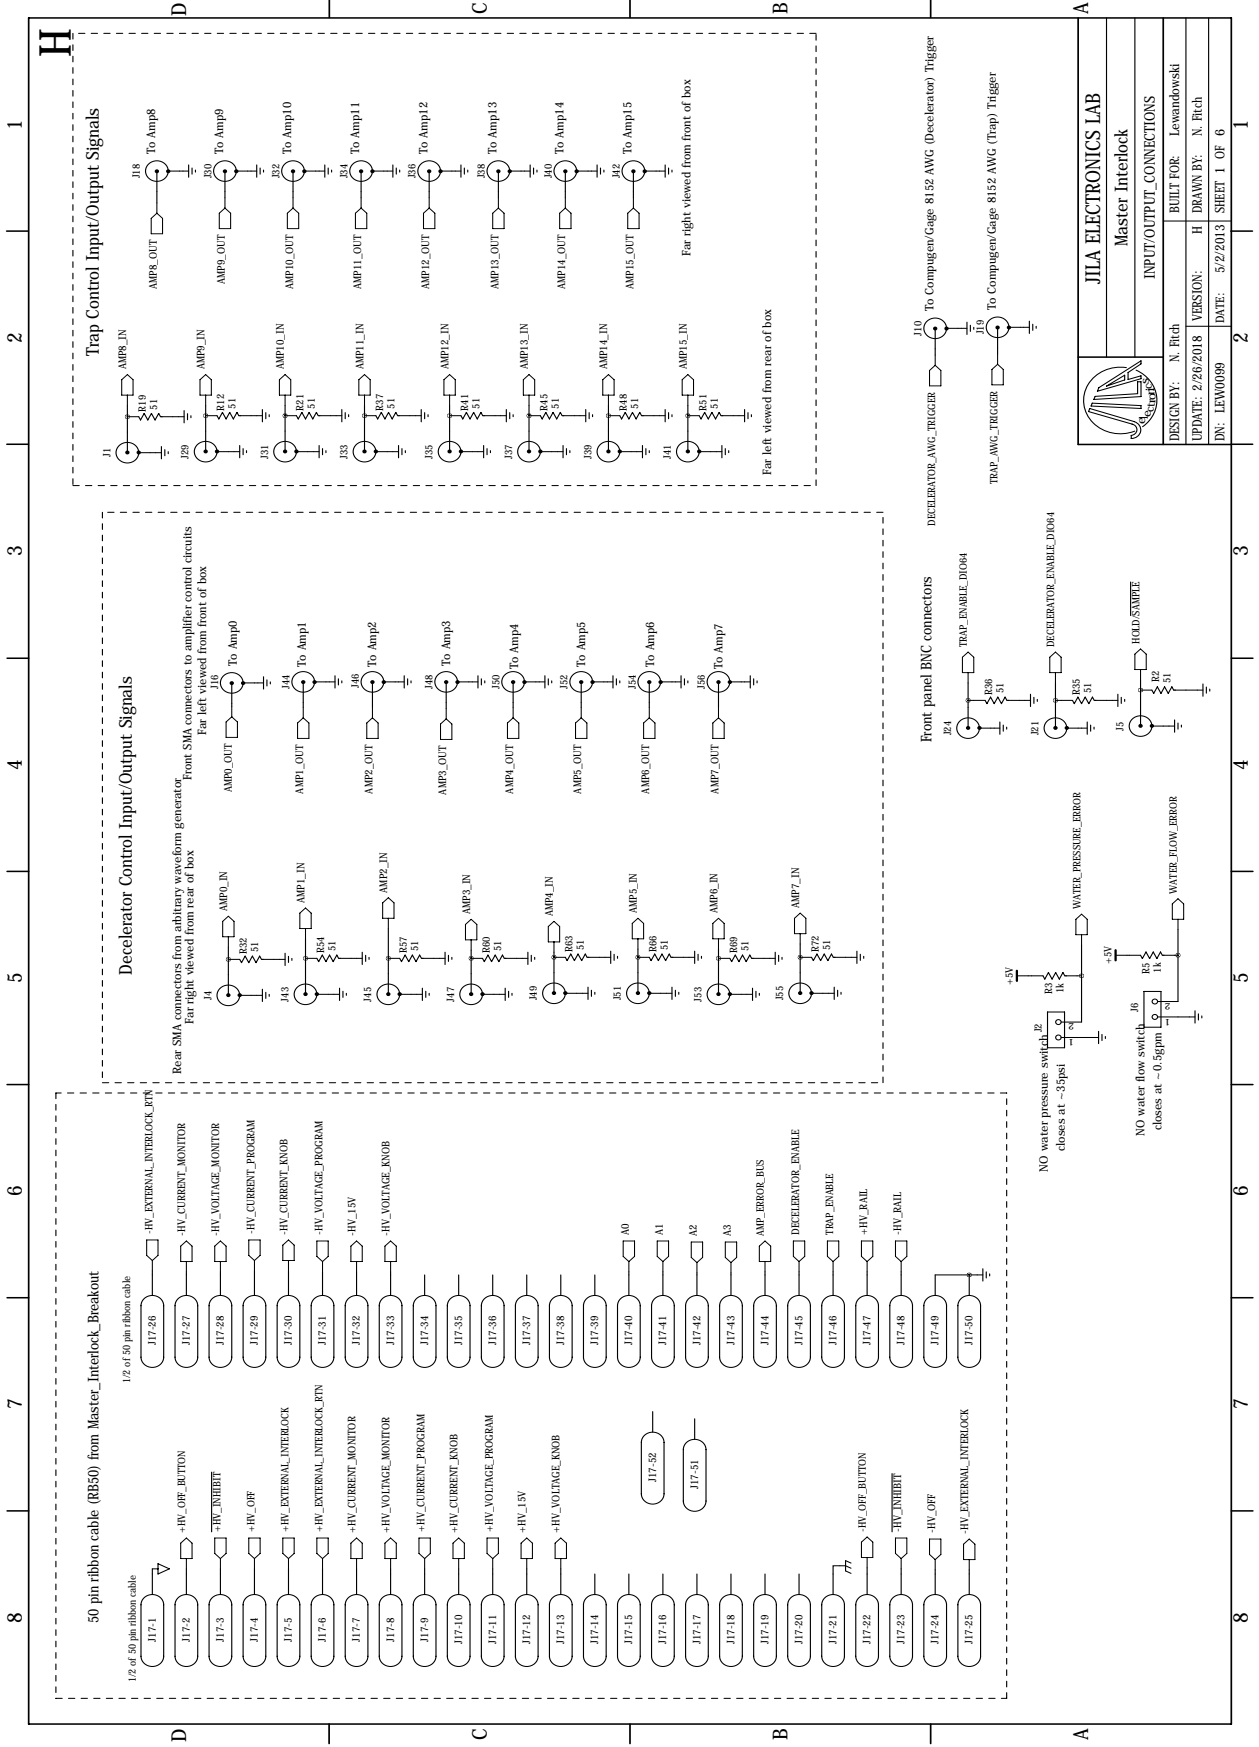
	JILA ELECTRONICS LAB	
	Outside Wall THERMAL_PADS	
DESIGN BY: N. Fitch	BUILT FOR: Lewandowski	
UPDATE: 12/13/2013	VERSION: H	DRAWN BY: N. Fitch
DN: <DN>	DATE: 5/2/2013	SHEET 1 OF 1






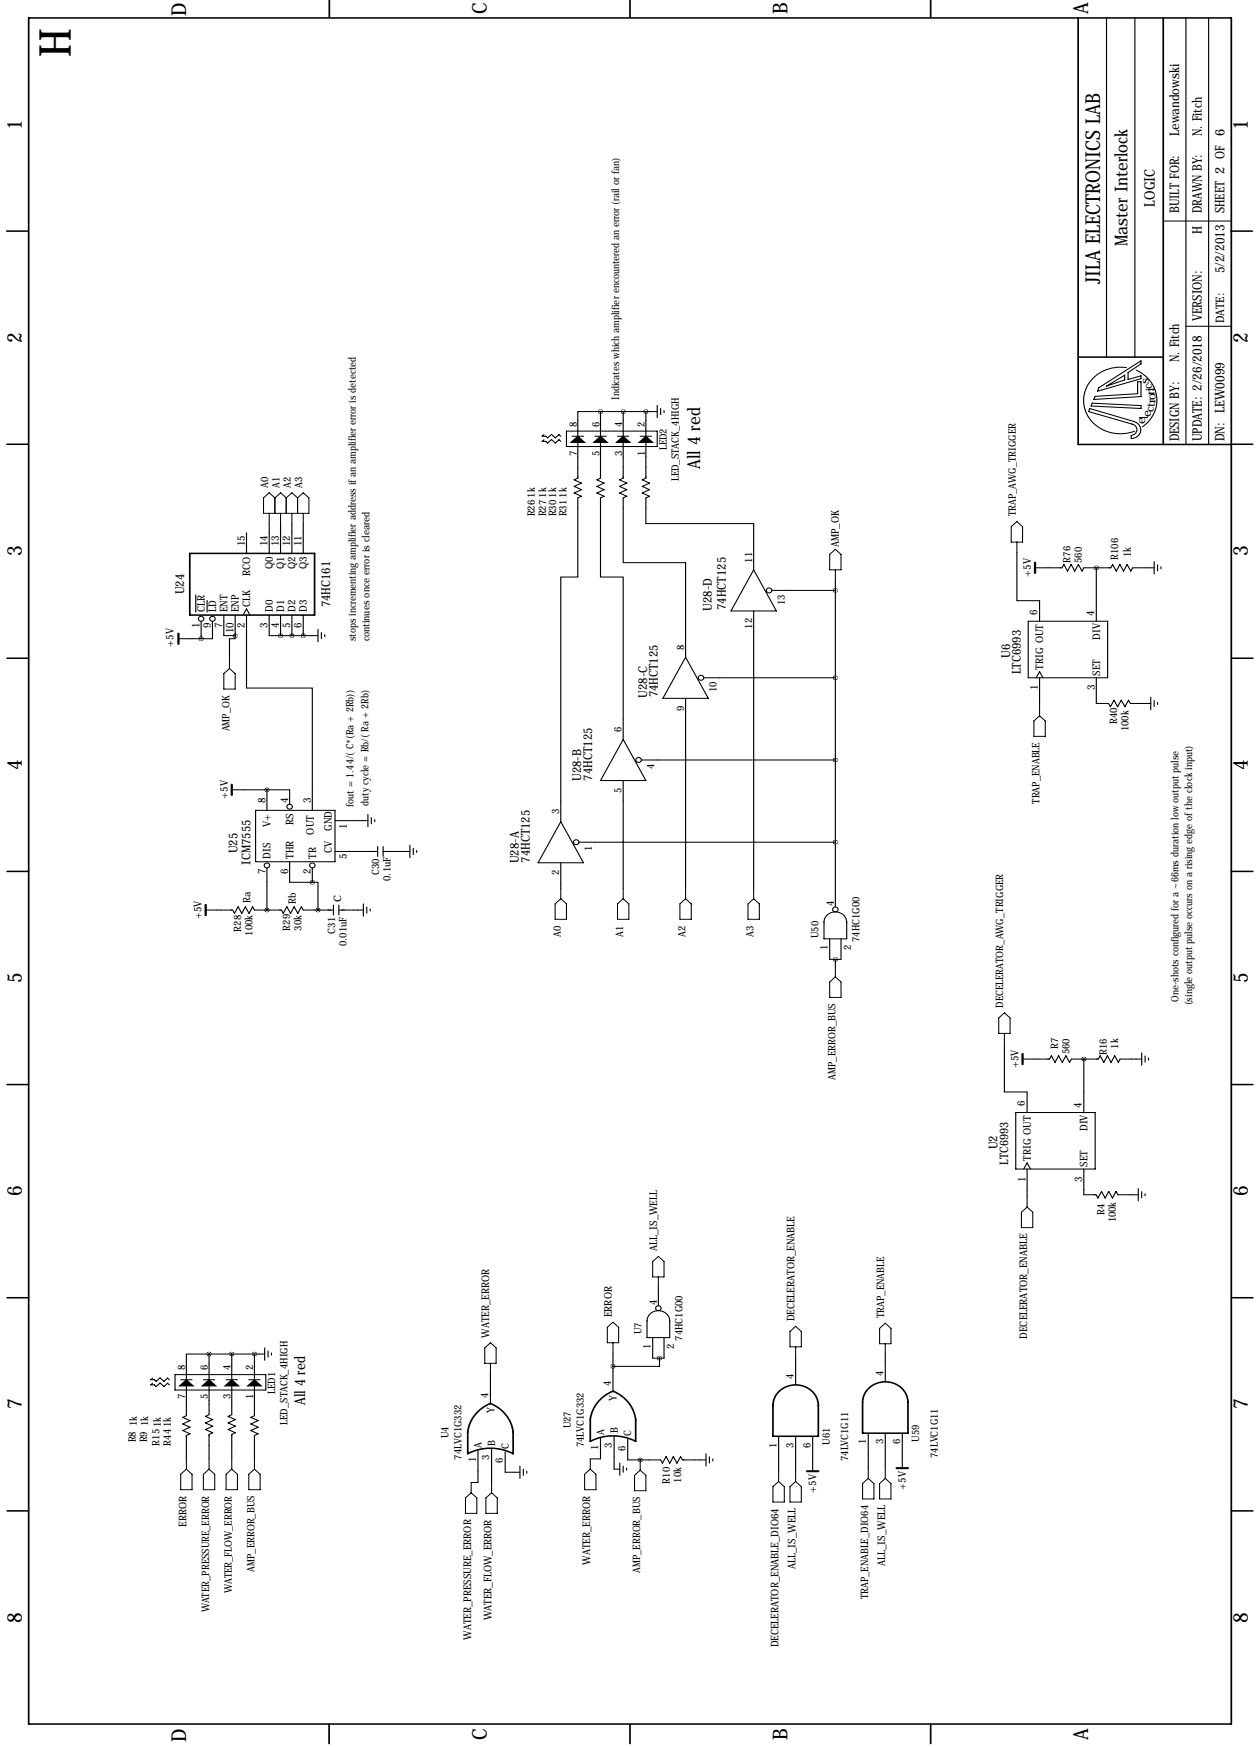
JILA ELECTRONICS LAB
Amplifier Breakout
CURRENT_AND_VOLTAGE_MONITORS

DESIGN BY: N. Fitch
 BUILT FOR: Lewandowski
 UPDATE: 2/26/2018
 VERSION: H
 DRAWN BY: N. Fitch
 DSN: LEW0089
 DATE: 5/3/2013
 SHEET 2 OF 2



JILA ELECTRONICS LAB
Master Interlock
INPUT/OUTPUT CONNECTIONS

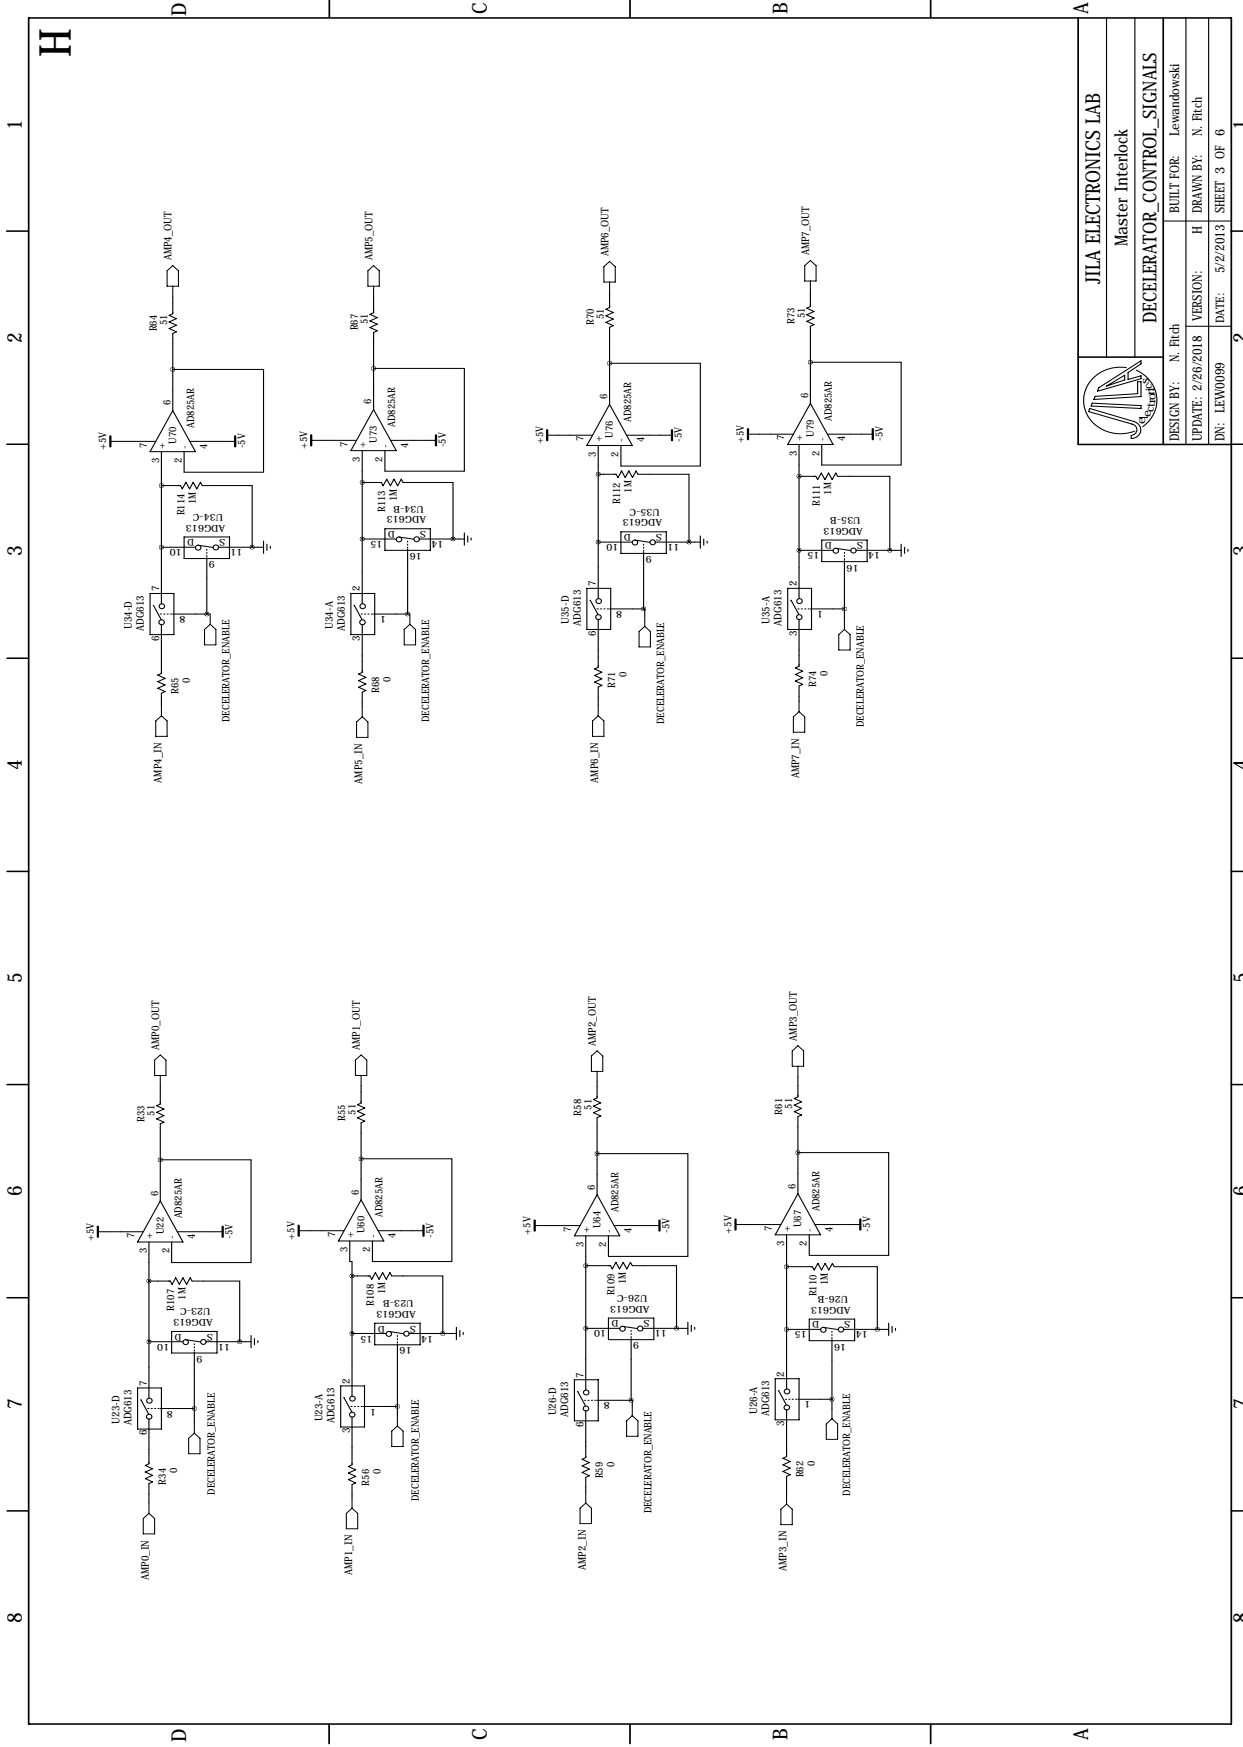
DESIGN BY: N. Rirch	BUILT FOR: Lewandowski
UPDATE: 2/26/2018	VERSION: H
DRAWN BY: N. Rirch	
DN: LEW00089	DATE: 5/2/2013
	SHEET 1 OF 6





JILA ELECTRONICS LAB
Master Interlock
LOGIC

DESIGN BY: N. Rich	BUILT FOR: Lewandowski
UPDATE: 2/26/2018	VERSION: H
DRAWN BY: N. Rich	
DATE: 5/2/2013	SHEET 2 OF 6

One-shots configured for a 40ns duration low output pulse (single output pulse occur on a rising edge of the clock input)

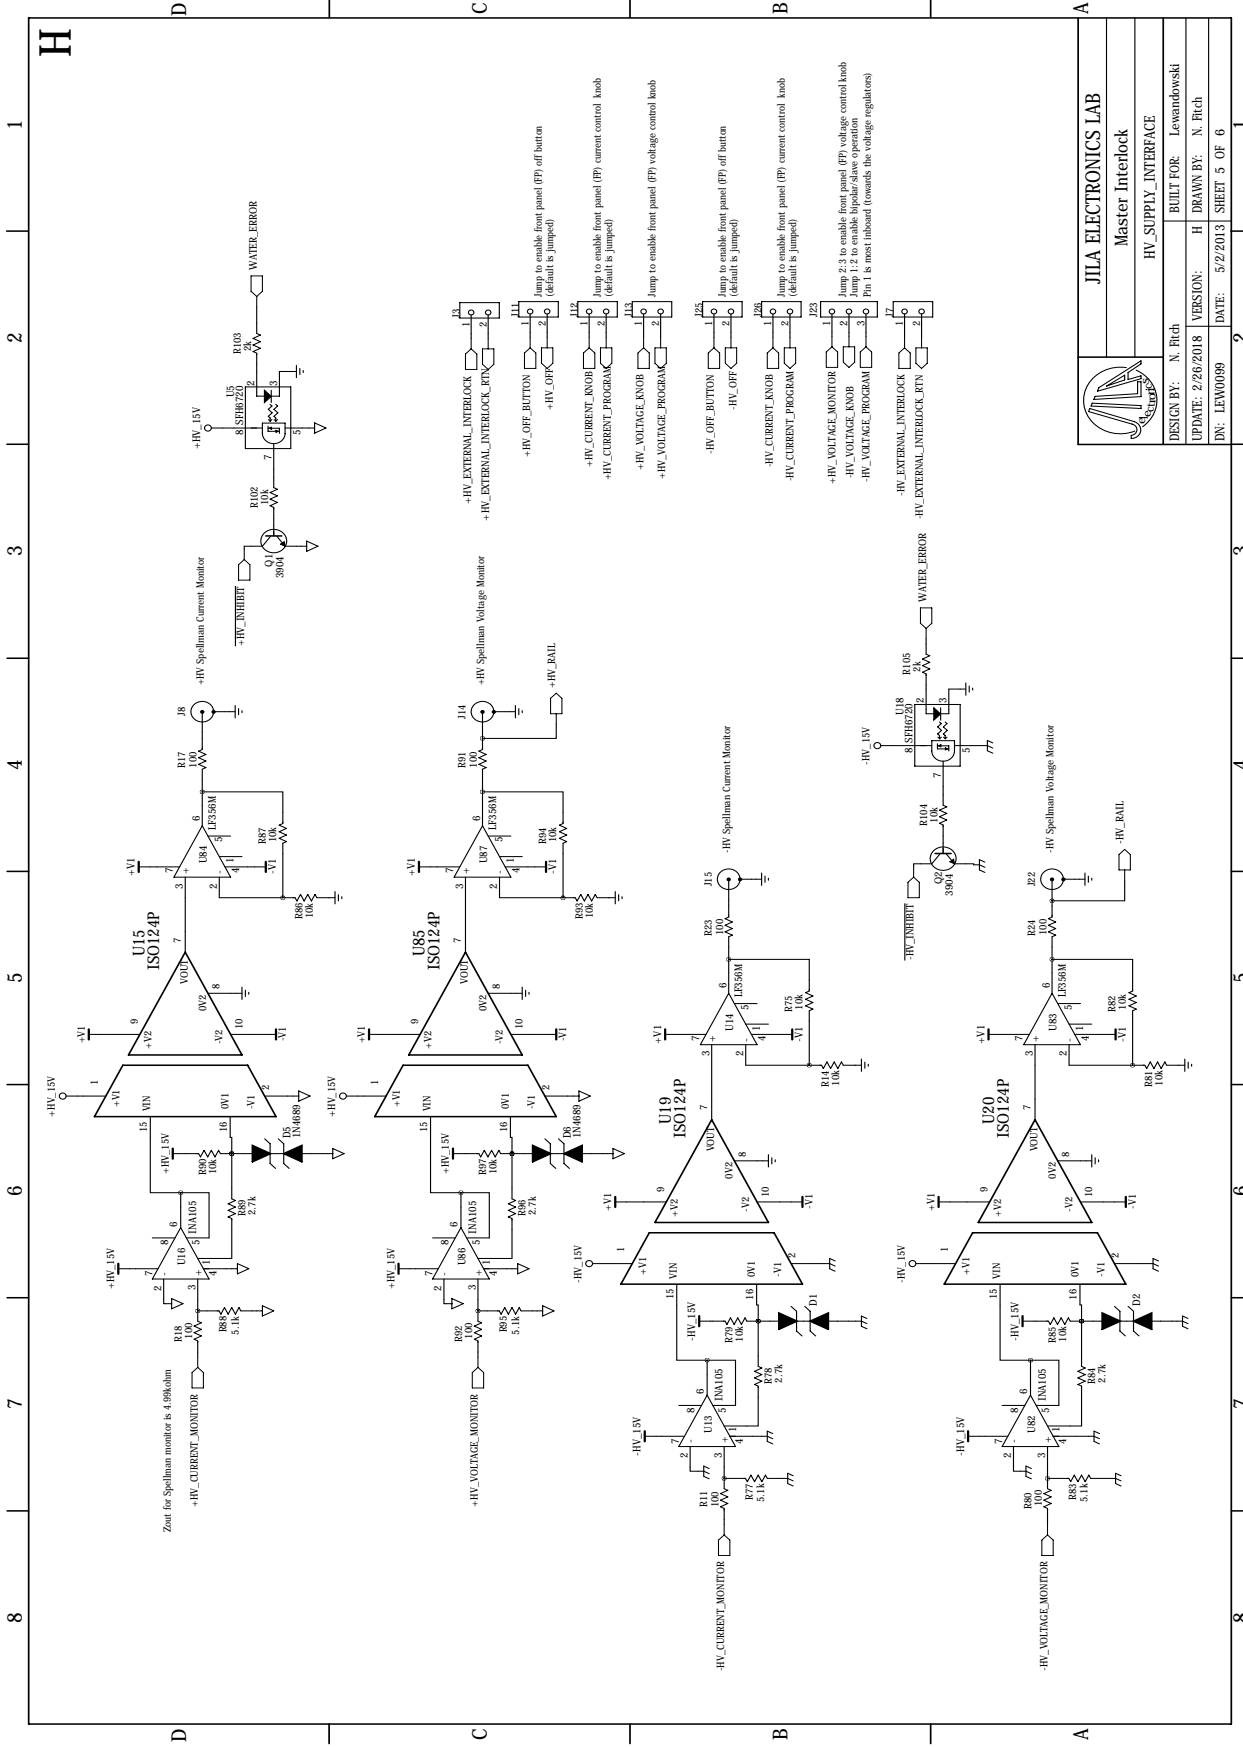


	
JILA ELECTRONICS LAB	
Master Interlock	
DECELERATOR CONTROL SIGNALS	
DESIGN BY: N. Rirch	BUILT FOR: Lewandowski
UPDATE: 2/26/2018	VERSION: H
DRAWN BY: N. Rirch	
DN: LEW00089	DATE: 5/2/2013
SHEET 3 OF 6	



JILA ELECTRONICS LAB
Master Interlock
TRAP_CONTROL SIGNALS

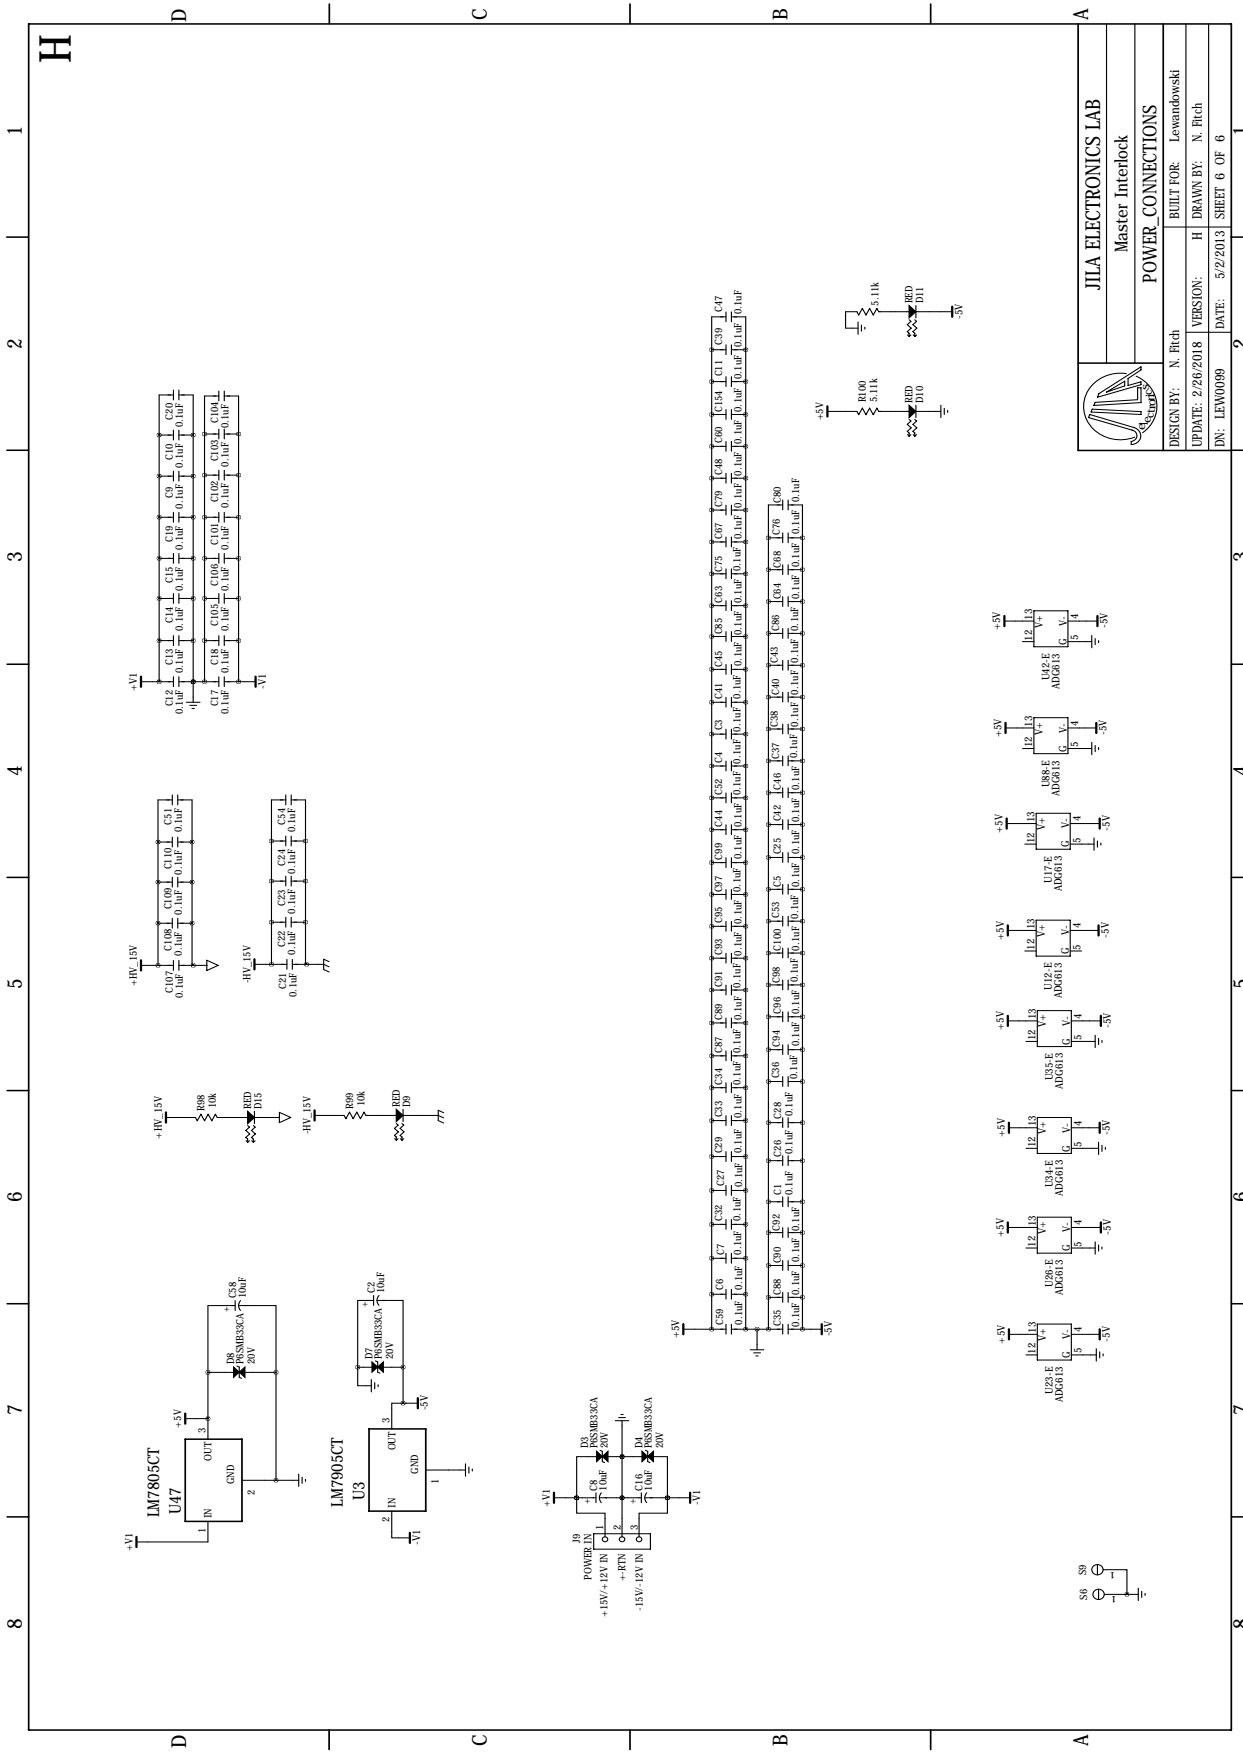
DESIGN BY: N. Rirch	BUILT FOR: Lewandowski	VERSION: H
UPDATE: 2/25/2018	DRAWN BY: N. Rirch	DATE: 5/2/2013
DN: LEW00089 SHEET 4 OF 6		



JILA ELECTRONICS LAB	
Master Interlock	
HV_SUPPLY_INTERFACE	
DESIGN BY: N. Ritch	BUILT FOR: Lewandowski
UPDATE: 2/26/2018	VERSION: H
DN: LEW0089	DATE: 5/2/2013
SHEET 5 OF 6	

1 2 3 4 5 6 7 8

8 7 6 5 4 3 2 1



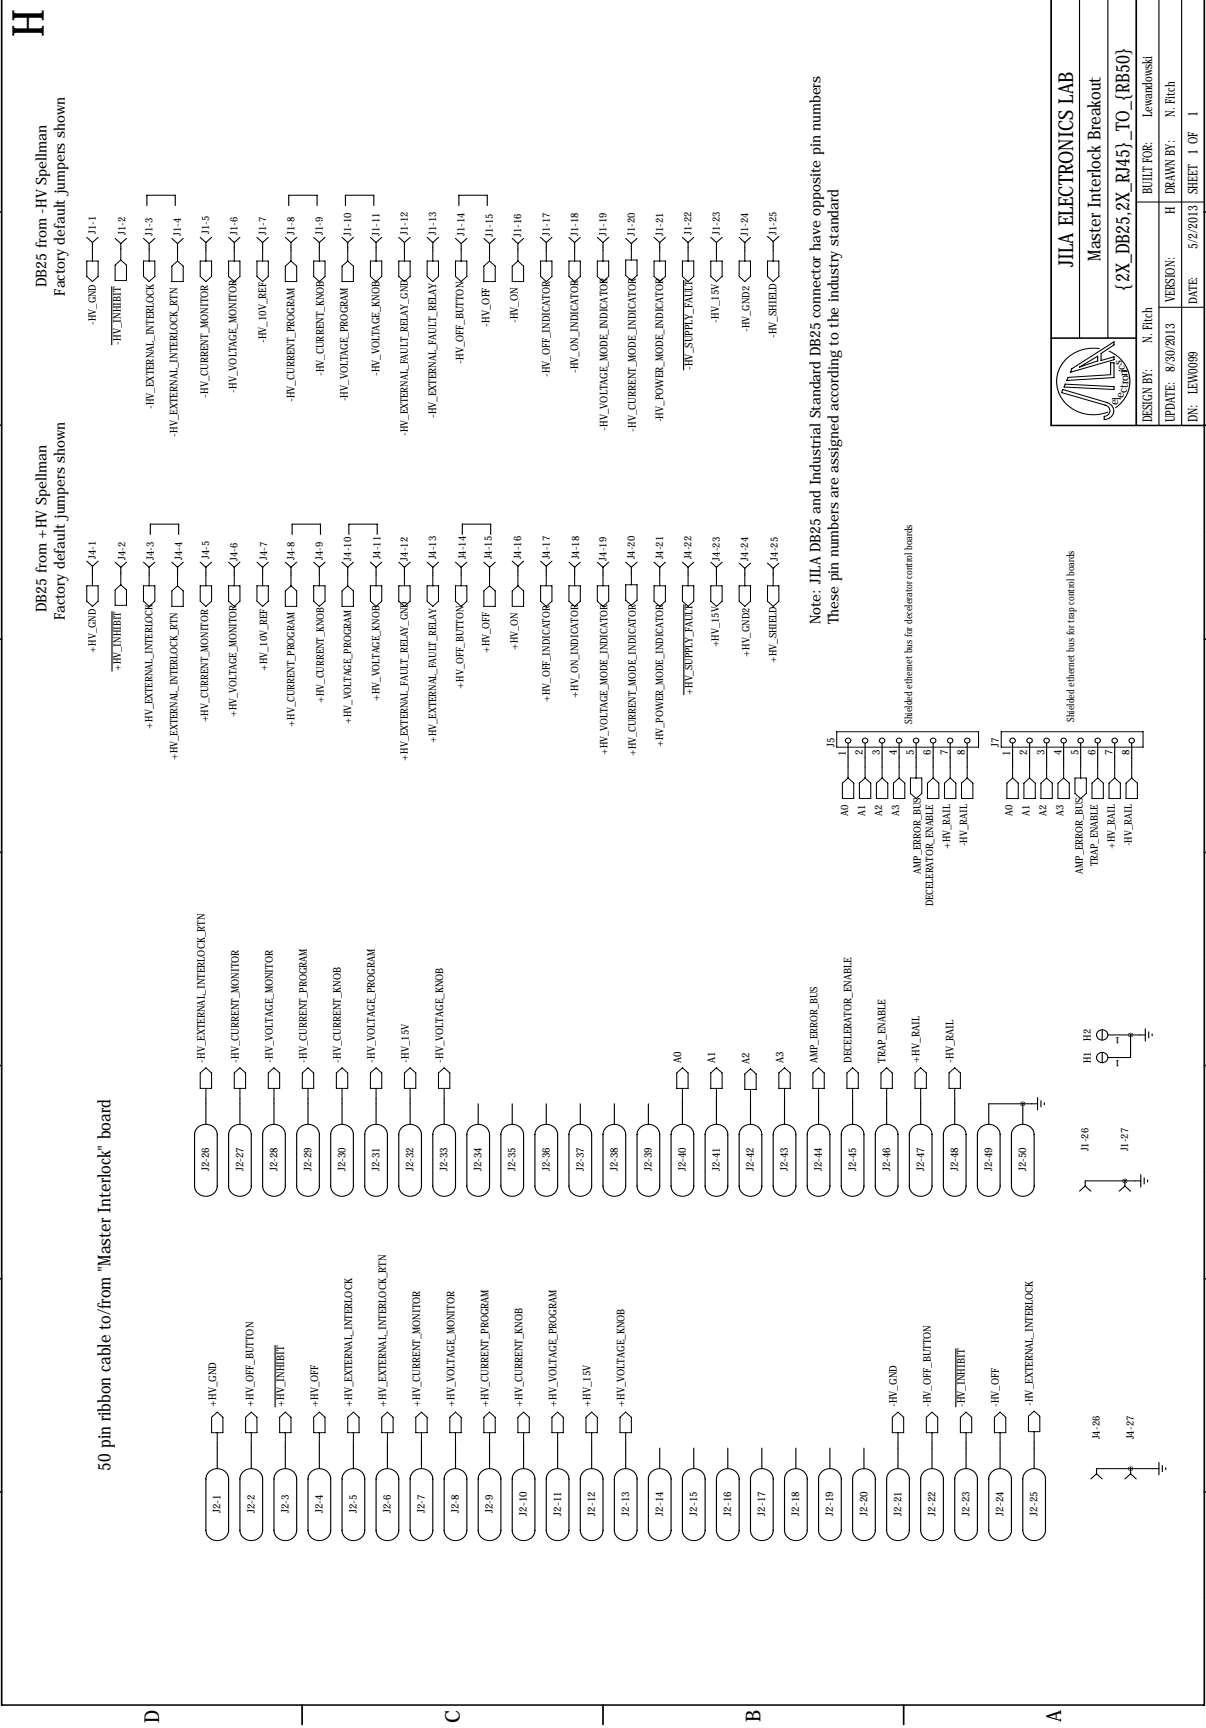
JILA ELECTRONICS LAB
Master Interlock
POWER CONNECTIONS

DESIGN BY: N. Ririh
UPDATE: 2/26/2018
VERSION: H
DRAWN BY: N. Ririh

DN: LEW00089
DATE: 5/2/2013
SHEET 6 OF 6

56	50	
58	50	
U17, U18, U19, U20, U21, U22	AD620	AD620
U3, U4	LM7905CT	LM7905CT
U47, U48	LM7805CT	LM7805CT

8 7 6 5 4 3 2 1



50 pin ribbon cable to/from "Master Interlock" board

Note: JILA DB25 and Industrial Standard DB25 connector have opposite pin numbers. These pin numbers are assigned according to the industry standard

JILA ELECTRONICS LAB Master Interlock Breakout {2X_DB25_2X_RJ45}_TO_{RB50}	
DESIGN BY: N. Fitch	BUILT FOR: Levandowski
UPDATE: 8/30/2013	VERSION: H
DATE: 5/2/2013	DRAWN BY: N. Fitch
DN: LEW009	SHEET 1 OF 1

Appendix B

Modifications summary

This appendix contains a concise summary of all the changes discussed in Chapter 5 sorted by board. The original values referenced here refer to the amplifier circuits presented in Noah Fitch's dissertation [95] and the update values are the ones given in Appendix A.

B.1 Amplifier bank

- The change to R21 is optional if Q8 or Q7 are not installed, see Chapter 5 for an explanation.
- The following list of parts do not have a part number listed, they are all MMSZ4701 15V zener diodes: D1, D12, D22, D23, D40, D52, D55, D67, D70, D82, D85, D97, D100, D112, D115, D127, D130, D142, D145.
- The brass shield on leader control circuit is no longer used.
- The voltage limiters remain on the circuit diagram but they are not installed. They are components: Q3-8, Q17-18, Q20, Q21, Q24.
- The bottom bank current monitors was shorted by connecting TP4 to TP7. This removes it from the circuit but it is left in for structural rigidity of the elevated ethernet board.

Component number	Original value	Updated value
R21	470 Ω	430 Ω
R159	300 Ω	100 Ω

B.2 Amplifier breakout boards

Component number	Original value	Updated value
C4	330 pF Ω	470 pF Ω
C6	180 pF Ω	100 pF Ω

B.3 Control boards

Component number	Original value	Updated value
R13	51 Ω	1M Ω
R35	200 k Ω	390 k Ω
R38	200 k Ω	390 k Ω
R166	100k Ω	10k Ω
R167	10k Ω	7.5k Ω
R168	5.11k Ω	7.5k Ω
R170	10k Ω	7.5k Ω
R171	100k Ω	10k Ω
R173	5.11k Ω	7.5k Ω

B.4 Master interlock

Pin 3 of the following components must be tied to ground using a 1 M Ω resistor: U8, U22, U29, U32, U33, U41, U45, U48, U54 U60, U64, U67, U70, U73, U76, U79

B.5 Insulation

In order to prevent current oscillations at output voltages greater than ± 8 kV, the amplifier banks and cables are potted with RTVS 3-95-1 and the high-voltage probes are potted with Hysol ES1001.