

INFORMATION TO USERS

This manuscript has been reproduced from the microfilm master. UMI films the text directly from the original or copy submitted. Thus, some thesis and dissertation copies are in typewriter face, while others may be from any type of computer printer.

The quality of this reproduction is dependent upon the quality of the copy submitted. Broken or indistinct print, colored or poor quality illustrations and photographs, print bleedthrough, substandard margins, and improper alignment can adversely affect reproduction.

In the unlikely event that the author did not send UMI a complete manuscript and there are missing pages, these will be noted. Also, if unauthorized copyright material had to be removed, a note will indicate the deletion.

Oversize materials (e.g., maps, drawings, charts) are reproduced by sectioning the original, beginning at the upper left-hand corner and continuing from left to right in equal sections with small overlaps. Each original is also photographed in one exposure and is included in reduced form at the back of the book.

Photographs included in the original manuscript have been reproduced xerographically in this copy. Higher quality 6" x 9" black and white photographic prints are available for any photographs or illustrations appearing in this copy for an additional charge. Contact UMI directly to order.

UMI

**A Bell & Howell Information Company
300 North Zeeb Road, Ann Arbor, MI 48106-1346 USA
313/761-4700 800/521-0600**

Demonstration of Trapped
Single Laser Cooled Indium Ions

by

Eric A. Burt

A dissertation submitted in partial fulfillment of
the requirements for the degree of

Doctor of Philosophy

University of Washington

1995

Approved by *David Nagay*
(Chairperson of Supervisory Committee)

Program Authorized
to Offer Degree DEPARTMENT OF PHYSICS

Date AUGUST 16, 1995

UMI Number: 9609597

UMI Microform 9609597

Copyright 1995, by UMI Company. All rights reserved.

**This microform edition is protected against unauthorized
copying under Title 17, United States Code.**

UMI

**300 North Zeeb Road
Ann Arbor, MI 48103**

In presenting this dissertation in partial fulfillment of the requirements for the Doctoral degree at the University of Washington, I agree that the Library shall make its copies freely available for inspection. I further agree that extensive copying of this dissertation is allowable only for scholarly purposes, consistent with "fair use" as prescribed in the U.S. Copyright Law. Requests for copying or reproduction of this dissertation may be referred to University Microfilms, 1490 Eisenhower Place, P.O. Box 975, Ann Arbor, MI 48106, to whom the author has granted "the right to reproduce and sell (a) copies of the manuscript in microform and/or (b) printed copies of the manuscript made from microform."

Signature _____

A handwritten signature in black ink, appearing to be a stylized 'E' followed by a large, sweeping flourish that ends in a horizontal line.

Date _____

8/16/95

University of Washington

Abstract

Demonstration of Trapped Single Laser Cooled Indium Ions

by Eric A. Burt

Chairperson of Supervisory Committee: *Professor Warren Nagourney*
Department of Physics

Single Indium ions have been confined in a RF Paul-Straubel trap, positioned within an ultra high vacuum environment. Indium is a group IIIA atom and, in its singly ionized state, has a $J = 0 - J = 0$ “clock” transition that is highly immune to shifts from ambient fields and field gradients. The immunity to field induced shifts, together with the tight confinement characteristic of trapped single ions, makes a single trapped laser cooled Indium ion an ideal candidate for high precision measurement and frequency standards. Indium is the first group IIIA ion to be trapped singly. The apparatus used, first to electrically detect large clouds of Indium ions, second to optically detect such clouds, and finally to optically detect single ions is described in detail. Since the “cooling” transition in In^+ is an intercombination line, laser cooling is not as strong as in other ion species that have been trapped. Therefore, successful trapping and cooling of In^+ required the use of a precooling technique for loading which is described. In addition, due to the weak cooling transition, it was essential to drastically reduce instrumental scattering background. A technique for suppressing the background and substantially increasing the signal to noise ratio, developed for this experiment, is also described. Finally, spectroscopy on the cooling transition of a single laser cooled Indium ion is given.

TABLE OF CONTENTS

List of Figures	v
List of Tables	viii
Chapter 1: Introduction	1
Chapter 2: Theoretical Aspects of Ion Trapping and Cooling	10
2.1 Trap Characteristics	10
2.2 Optimal Trap Geometry	13
2.3 Trapping Ion Clouds	14
2.3.1 Buffer Gas Cooling	14
2.3.2 Heating Effects	16
2.3.3 Number of Ions Trapped	17
2.4 Trapping Single Ions	18
2.4.1 Laser Cooling	18
2.4.2 Optimal Detuning for Laser Cooling	21
2.4.3 Lamb-Dicke Confinement	21
2.4.4 Heating Effects	22
2.5 Expected Fluorescence Signal Strength	22
2.6 Saturation Power and Power Broadening	24
2.7 Quantum Jumps	25
2.8 Ambient Magnetic Fields and Optical Pumping	27
2.9 Hanle Effect Rotations and Signal Strength	27
Chapter 3: Indium	31
3.1 Group IIIA Ions	31
3.2 Cooling Transition Considerations: Choosing In^+	32
3.2.1 Transition Lifetimes	34

3.2.2	Expected Fluorescence Signal From Trapped Indium Ions	34
3.2.3	Laser Power Requirements	35
3.3	Hyperfine Structure and Zeeman Shifts	35
3.4	The Stark Shift	38
3.5	Isotopes	39
Chapter 4: Experimental Apparatus		40
4.1	Overview	40
4.2	Laser Sources	42
4.2.1	Dye Change Procedure	42
4.2.2	Reacquisition of Lasing after Dye Change	43
4.2.3	Jet Assembly Problems after Dye Change	44
4.3	Laser Stabilization	45
4.4	Short Term Stabilization	47
4.4.1	Apparatus	47
4.4.2	Electronics	50
4.4.3	Servo Loop Stability	55
4.5	Drift Stabilization and Location of Indium Transitions: Tellurium Sat- uration Cell	56
4.5.1	Apparatus	57
4.5.2	Electronics	59
4.5.3	Operation of Laser Stabilization Apparatus	59
4.6	Second Harmonic Generation	63
4.6.1	BBO Crystal	63
4.6.2	SHG Enhancement Cavity	68
4.6.3	Cavity Optics	69
4.6.4	Cavity Locking Electronics	73
4.7	Beam Shaping and Input Optics	74
4.8	Vacuum Chamber	77
4.8.1	Cleaning and Bakeout Procedure	80
4.8.2	Buffer Gas Plumbing	83
4.9	Trapping Apparatus	85

4.9.1	Trapping Electronics	92
4.10	Detection Strategy	93
4.10.1	Electrical Detection	94
4.10.2	Optical Detection	96
4.11	Background	100
Chapter 5:	Background Suppression	102
5.1	Overview of the Background Suppression Method	103
5.2	Theoretical Analysis of a Two State System Response to Square Wave Modulated Excitation	104
5.3	Background Suppression Apparatus	113
5.4	Background Suppression Results	116
Chapter 6:	Results	120
6.1	Buffer Gas Cooled Ion Clouds	120
6.1.1	Loading Procedure	121
6.1.2	Electrical Detection	122
6.1.3	Optical Alignment Procedures	123
6.1.4	Optical Fluorescence Detection	125
6.1.5	Tellurium Saturation Spectroscopy	126
6.2	Laser Cooled Single Indium Ions	136
6.2.1	Single Ion Search Strategy	137
6.2.2	Single Ion Search Space	139
6.2.3	Chronology of Single Ion Search	140
6.2.4	Loading and Observing Laser Cooled In^+ ions	147
6.2.5	Micromotion	154
6.2.6	Determination of the Correct In^+ Cooling Transition Frequency Relative to a Tellurium Line.	156
6.2.7	Single Ions?	156
6.2.8	Loading Parameters	157
6.2.9	Alignment Procedure	158
6.2.10	Loading Procedure	160

6.3	Conclusions	161
6.4	Where to Next?	162
	Bibliography	164
	Appendix A: Indium Ion Summary	170
	Appendix B: Some Useful Formulae	173
B.1	Correction Plate Displacement	173
B.2	Gaussian Fits	175
B.3	Summary of Useful Equations	176
	Appendix C: Electrical Circuits	178
C.1	Trap RF Drive Amplifier	178
C.2	Electrical Detection Preamp	178
C.3	Short Term Laser Stabilization Electronics	178
C.4	Drift Stabilization Electronics	181
C.5	Enhancement Cavity Locking Electronics	188
C.6	Gating and Detection Electronics	188

LIST OF FIGURES

3.1	Indium II Level Diagram	33
4.1	Experimental Apparatus	41
4.2	Short Term Stabilization Apparatus	47
4.3	Tellurium Saturation Cell	56
4.4	BBO Crystal Faces	66
4.5	Second Harmonic Generation Enhancement Cavity	70
4.6	Vacuum Chamber	78
4.7	Helium Leak Apparatus	84
4.8	Paul Trap	86
4.9	Paul-Straubel Trap	90
4.10	Ion Trapping Electronics Block Diagram	92
4.11	Electrical Detection Block Diagram	95
4.12	Optical Detection Apparatus	96
5.1	Light Modulation And Detection Gating	105
5.2	Simulated Excited State Population With Modulated Excitation	107
5.3	Simulated Data For Modulated Fluorescence Vs. Modulation Period	109
5.4	Fluorescence Vs. Modulation Period For Several Duty Cycles And Random Laser Noise	110
5.5	Simulated Data For Excited State Population Vs. Time For Two Mod- ulation Periods	111
5.6	Simulated Data For Modulated Fluorescence as a Fraction of Unmod- ulated Fluorescence	112
5.7	Fluorescence Vs. Detection Time	114
5.8	Simulated Excited State Population Vs. Laser-Off Time With 10 MHz Laser	115
5.9	Strong Fluorescence Signal, With and Without Background Suppression	117

5.10	Comparison Of A Weak Ion Cloud Signal With, And Without Background Suppression	118
5.11	Weak Single Ion Signal, With And Without Background Suppression	119
6.1	Fluorescence From The $In^+ \ ^1S_0 - ^3P_1$ Transition For An Ion Cloud	127
6.2	Ion Cloud Signal With SHG Enhancement And Background Suppression	128
6.3	Tellurium Mapping Apparatus	130
6.4	Dual Scan Of $In^+ \ 9/2 - 7/2$ Line And Tellurium Saturation Signal	132
6.5	Data Fit For Tellurium Map	134
6.6	In^+ Ion $9/2 - 11/2$ HF And Tellurium Spectra Taken Simultaneously	135
6.7	Vacuum Ion Cloud Signal	142
6.8	First Laser Cooled In^+ Ion Signal	148
6.9	First Laser Cooled Ion Signal With a Narrower Scan	149
6.10	Comparison Of A Laser Cooled Ion Signal With, And Without Background Suppression	151
6.11	Best Laser Cooled Ion Signal	152
6.12	Spectrum Of The Narrowest Laser Cooled Ion Signal Obtained	153
6.13	The Effect Of Heating The Trap On A Laser Cooled Ion Signal	155
B.1	Astigmatism Correction Plate Geometry	174
C.1	RF Drive Amplifier	179
C.2	Electrical Detection Preamp	179
C.3	Detection Electronics For Laser Stabilization	182
C.4	Oscillator And Mixer Electronics For Laser Stabilization	183
C.5	Integration Electronics For Laser Stabilization	184
C.6	Slow Path Transfer Electronics For Laser Stabilization	185
C.7	Redundant Locking Electronics For Laser Stabilization	186
C.8	Drift Stabilization Electronics Block Diagram	187
C.9	SHG Locking Electronics Schematic—Part 1	189
C.10	SHG Locking Electronics Schematic—Part 2	190
C.11	SHG Locking Electronics Schematic—Part 3	191
C.12	Background Suppression Gating Electronics—Part 1	192

C.13 Background Suppression Gating Electronics—Part 2	193
C.14 Background Suppression Gating Electronics—Part 3	194

LIST OF TABLES

4.1	Brewster Cut BBO Parameters	67
4.2	Paul-Straubel Trap Parameters	89
6.1	Tellurium Map Sideband Fit Data	133
A.1	Indium II Frequencies And Lifetimes	170
A.2	Isotope Data	171
A.3	Laser Saturation Power For The Indium II Transitions	171
A.4	Estimated Fluorescence Signal	172

ACKNOWLEDGMENTS

The work described in this dissertation was supported by a grant from the Air Force Office of Scientific Research. I would like to thank my advisor, Prof. Warren Nagourney, for his guidance on this difficult experiment. I was also fortunate to have access to the large amount of experience in the field of single ion trapping represented in Hans Dehmelt's lab where I worked. Many discussions with Prof. Nan Yu and Dr. Xin Xin Zhao in addition to Prof. Nagourney formed the foundation of my understanding of the many aspects of the Indium experiment. Prof. Norval Fortson was the first member of the University of Washington Physics Department faculty that I met. His guidance during my entire time as a graduate student was invaluable. In addition, Prof. Protik Majumder and Prof. Steve Lamoreaux were always willing to discuss any physics topic that I brought up. Rarely does a physics experiment succeed without the assistance of several physics department support personnel. The Indium experiment was no exception. Ron Musgrave in the machine shop was always willing to patiently explain machining techniques and help with the multitude of machining projects that have to be carried out in order to set up an optical table. Our glass blower, Bob Morely, fabricated our quartz vacuum chamber and responded quickly and professionally to requests for modifications. Finally, Phil Williams and Dr. Bryan Venema in the electronics shop, were always willing to take questions and to help solve electronics problems. There is probably no piece of electronics in our experiment that Phil didn't have a part in to some extent. I would also like to thank my thesis reading committee of Prof. Nagourney, Prof. Norval Fortson, and Prof. Steve Lamoreaux for their careful reading of this manuscript and their many useful comments. My wife Peggy's proofreading skills were also invaluable in this regard, not to mention her unflagging support over the 6 years that I was

in the Ph.D. program. The consistent encouragement of my parents, George Burt and Dierdre Burt Anderson is particularly appreciated. Without this, and the financial support of my mother I would not have been able to start down the physics path.

Chapter 1

INTRODUCTION

In 1980 Hans Dehmelt and collaborators [Neu80] first successfully trapped a single Ba^+ ion. This tremendous feat was the realization of one of the main steps towards Dehmelt's goal of isolating a single atomic particle at rest in space: a "mono-ion oscillator". Such a particle would be free from outside perturbations and observable for long periods of time opening the way to measurements of unprecedented precision. Since the first single ion was trapped, much work has been done in this field and several other ions have been successfully trapped. Until recently, all of these ions were Alkali-like, having non zero total angular momentum in the transition which is studied. The angular momentum causes splittings and shifts in the transition due to couplings with external fields and hence limits the precision of measurements that can be made. Singly ionized Group IIIA elements do not have this drawback, but the required light sources have traditionally been difficult to synthesize. Since deciding on an ion to study is largely dictated by the availability of the necessary light sources, Group IIIA ions had not been attempted. It has been the goal of this thesis project to extend the field of single ion trapping into the realm of the Group IIIA ions by trapping, cooling, and performing high resolution spectroscopy on single Indium ions.

The story of ion trapping has its beginnings in the 1950's with Wolfgang Paul's experiments on radio-frequency (RF) traps [Pau58]. It is well known that a 3 dimensional trapping potential cannot be created from a single DC electric or magnetic field. A Penning trap consists of DC potential applied between a ring and two endcaps *and* a magnetic field along the axis of symmetry. While this creates a good trapping potential, the presence of the magnetic field may be undesirable. A single field can provide a local energy minimum for an ion if it is inhomogeneous and it is allowed to oscillate. An RF potential applied between a ring electrode and endcaps (using the same geometry as in the Penning configuration) creates this energy minimum. The effective potential created by the inhomogeneous RF field is usually referred to as the

psuedopotential. These RF quadrupole traps are referred to now as Paul traps.

High resolution spectroscopy on ion clouds in an RF trap was first accomplished in 1966 by Fortson, Major and Dehmelt [For66]. This experiment studied hyperfine transitions in the ground state of ${}^3\text{He}^+$ which are in the microwave region of the spectrum. Shortly after that, in 1967 Dehmelt [Deh67] experimented with large clouds of ions cooled with a Helium buffer gas in a Paul trap. The ions were observed electrically by driving the cloud coherently and detecting an induced current in an external tank circuit. In this same paper Dehmelt gave a detailed analysis of trap parameters such as expected trap depth, frequency, and capacity, as well as expected cooling and heating effects.

In 1973 Dehmelt put forward the idea of a group IIIA mono-ion oscillator [Deh73]. This proposal was important for several reasons. Because of its high degree of isolation and confinement, a single trapped ion has a very good potential for high precision work. The 1973 proposal is the first to consider this. Also it suggests the use of Tl^+ , which is a group IIIA ion, as the ion to be studied. Group IIIA ions have a $J = 0 - J = 0$ “field independent” transition which is largely immune to shifts from external magnetic fields and electric field gradients making them excellent candidates for high resolution spectroscopy. This proposal was expanded on in 1975 [Deh75, Win75]. The most notable addition was the suggestion that the very weak $J = 0 - J = 0$ transition could be observed by the absence of fluorescence in a strong transition sharing the same ground state. In this scheme both transitions would be excited with the atom going into the metastable state infrequently. When it did, fluorescence from the strong transition would be shut off. Dehmelt referred to this as “electron shelving” and noted that the technique could serve as a million-fold amplifier for detection of the weak transition to the metastable state. This is the first mention of the possibility of observing quantum jumps—something that there was much debate over and wouldn’t be observed experimentally for another 11 years.

An essential part of being able perform high resolution spectroscopy on single ions (indeed, perhaps even observing them) is being able to cool them to the point where the Doppler width of the transition being studied is smaller than its natural width. In 1975 Hansch and Schawlow [Han75] and Wineland and Dehmelt [Win75] first proposed the use of lasers to perform this cooling. Ashkin [Ash70a, Ash70b, Ash78] and Letokhov [Let77] also discuss laser cooling in different contexts. In this technique

the laser is detuned below (“to the red of”) the cooling transition as it excites the ion. Because of its motion and owing to the Doppler effect, the ion is able to absorb energy from the laser as it travels towards the laser beam. The absorption process can be thought of as a collision between the ion and an incoming photon which causes the ion to slow down. When the excited ion re-emits a photon it receives a recoil kick, but these are directed isotropically (actually a dipole pattern) and the average kick after many absorption-emission cycles is very small. Consequently the time-averaged result is that the ion emits more energy than it absorbs and in this way is cooled.

Laser cooling of trapped ions was first observed by Wineland et al. [Win78] in a small cloud of Mg^+ ions and concurrently by Neuhauser et al. [Neu78] in a small cloud of Ba^+ ions. In the latter case, the cloud was first loaded in the presence of a buffer gas. Then the gas was pumped out while the laser was detuned 500 MHz. The number of ions could roughly be determined by the size of the cloud (down to the diffraction limit) and occasionally cloud sizes small enough to correspond to 1 ion were observed.

There are two important regimes of laser cooling in ions: Doppler cooling and sideband cooling. Because of the oscillatory motion of the ion in the trap, the fluorescence given off after it has been excited will have sidebands spaced at the trap frequency or secular frequency. When the ion is hot, the Doppler width is usually much broader than the sideband spacing and the sidebands all coalesce to form the usual Doppler profile. However, even when the ion is cold, if the natural width is still larger than the sideband spacing, the sidebands will not be (clearly) resolved in the resulting fluorescence. This is called the Doppler cooling regime. The optimal laser detuning for cooling in this regime is one half the larger of the Doppler or natural width for the transition. Since the Doppler width will change as the ion is cooled, the laser detuning will have to be changed if cooling efficiency is to be maintained. The Doppler cooling limit is the temperature associated with one half of the transition’s natural width. On the other hand, if the natural width is smaller than the sideband spacing then the sidebands are resolved. In this case, the optimal laser detuning is the sideband spacing. The distinction between these two cooling regimes was first made by Wineland and Dehmelt early on [Win75] and analyzed by Dehmelt using a semi-classical argument in 1982 [Deh82]. A full quantum mechanical treatment was given by Wineland and Itano in 1979 [Win79a].

As far back as 1966 Dehmelt mentions the possibility of using trapped ions as a frequency standard and he is generally credited as being the first to have the idea. Building such a standard usually means locking a laser (in the case of an optical transition) to the transition. There are many problems that must be solved before such a standard can be converted into an atomic clock. Among these is that if an optical transition is being used there are no electronics fast enough to count the number of oscillations per second that are associated with this transition. Some sort of frequency down conversion technique into the microwave region is needed. Such a technique was proposed (and later developed) in 1979 by Wineland [Win79b] using the (relatively) slow cyclotron frequency of an electron in a Penning trap as the conversion mechanism. In this scheme "...the electron-cyclotron orbit is stabilized by balancing the synchrotron radiation...with power supplied by the focused laser beam... For the phase-lock condition, the measured cyclotron frequency is an exact submultiple...of the laser frequency." The laser is locked to the trapped ion and its absolute frequency is determined by measuring the cyclotron frequency of an electron in a Penning trap phase-locked to the laser.

As has already been mentioned, the first laser cooled single ion signal was seen in 1980 and hints of such a signal were seen as early as 1978. It is interesting that as recently as 1977, textbooks were claiming that while one could analyze the interaction of light with one atom (or ion), one would never be able to observe it directly [Cor77, p. 494]. The belief was that one would only be able see ensembles of atoms (ions) in actual experiments. There were indications of single ions as early as 1978 (the cloud size argument already mentioned) but in 1980 these initial signs were confirmed by seeing steps in the fluorescence during the loading process. The steps allowed the experimentalists to see each ion as it is loaded. This made it possible to determine exactly the amount of fluorescence associated with each ion and, therefore, when only one ion was present. The work was done with a Ba^+ ion. While the successful trapping and laser cooling of a single ion was a major milestone in the ion trapping story, it was by no means the end of the story. It still remained to cool the ion into the zero-point energy level of the trapping potential well. This would then represent the closest that one could get to an ion at rest. It should be noted that precision work does not necessarily require zero-point confinement. Rather, the elimination of the first order Doppler shift is the key requirement. Nevertheless, being able to

study an ion in the zero-point energy level of a quantum mechanical oscillator is interesting from a theoretical standpoint. Shortly after the trapping of the single Ba^+ ion, in 1981, Wineland and Itano observed a single laser cooled Mg^+ ion and measured its temperature at roughly 50 mK [Win81]. In 1983 Nagourney, Janik and Dehmelt [Nag83] lowered this temperature measurement for a single laser cooled Mg^+ ion to 5 mK.

Laser cooling of single ions had now been clearly demonstrated, but the ions that were worked with all had natural widths large compared to the motional sideband spacing so that the sidebands weren't resolved. This meant that the cooling limits reached were well above the trapping potential zero-point as well. One way to get around this limitation that is sometimes possible is to use a 2-photon transition which can be narrowed by detuning the two lasers from the two transitions involved. Steps towards this goal were taken by the Seattle group (Janik, Nagourney and Dehmelt) in 1985 who were able to resolve a 2-photon peak in Ba^+ [Jan85]. In this experiment they observed Doppler free spectroscopy of a single Ba^+ ion and were the first to demonstrate Lamb-Dicke confinement in optical spectroscopy [Dic53]. Lamb-Dicke confinement refers to the situation where the amplitude of oscillatory motion of the ion in the trap is much less than the wavelength of the radiation being used to excite it (hence "confinement"). The importance of Lamb-Dicke confinement is first, that the first order Doppler shift is eliminated and second, that the ion now exists in a very small region of space. The latter makes measurements insensitive to spatial inhomogeneities in the measuring device. Together these make possible high precision measurement.

As has already been mentioned, in 1975 Dehmelt suggested the use of "electron shelving" into a metastable state together with a strong transition sharing the same ground state, as an amplifier for transitions into the metastable state. The process of the strong fluorescence shutting off due to the shelving has come to be known as quantum jumps. Suddenly, in 1985, there was much discussion on quantum jumps and in 1985 and 1986 several groups predicted that this phenomena should be observable [Coo85, Coh85, Jav86, Sch86]. Shortly after that in 1986, three groups observed quantum jumps in rapid succession. The first was the Seattle group of Nagourney, Sandberg and Dehmelt [Nag86] who made the observation on Ba^+ . Next came the (then) National Bureau of Standards (NBS) group of Bergquist, Hulet, Itano, and

Wineland [Ber86] who made the observation in Hg^+ followed immediately by Sauter, Neuhauser, Blatt and Toschek [Sau86] who also made the observation in Ba^+ . These observations were very important because they opened the way to performing spectroscopy on these long lived (on the order of seconds) transitions which otherwise would be too weak to detect. With quantum jumps, spectroscopy on a metastable state sharing the same ground state as a strong transition is accomplished by counting the number of jumps that take place as a function of laser detuning. The laser exciting the strong transition, which is usually the cooling laser as well, is held detuned to the red for cooling purposes. The laser exciting the metastable state is then scanned across its transition and a signal averager counts the number of dark periods in each frequency channel. This number, as a function of frequency, is the spectrum.

Now that very narrow transitions could be observed, the next goal was to make the ions as cold as possible. For very cold ions the quantum energy levels of the trapping potential become apparent. The lowest energy level, known as the zero-point energy, is the ultimate target for laser cooling of ions. In 1989 the NBS group determined that they had cooled a single Hg^+ ion to the extent that it was in the zero-point energy level 95% of the time [Die89]. When the quantum number of the ion's energy level is low (on the order of 10), the quantum nature of the potential well makes itself apparent in different absorption strengths of the two motional sidebands. If n is the energy quantum number then the upper sideband strength is proportional to $\langle n \rangle + 1$ whereas the lower sideband strength is only proportional to $\langle n \rangle$ where the brackets indicate averaging. Of course if $\langle n \rangle$ is large, the difference between the two strengths is negligible. Only when $\langle n \rangle$ is quite low does the difference become observable. If one measures the ratio of the two strengths then $\langle n \rangle$ can be determined.

With the zero-point essentially reached, some groups turned their attention to extending these techniques to other trapping geometries and new ion species. One drawback of the Paul RF trap is that, while it is possible to isolate and cool a single ion into the zero-point energy, it is not possible to do so with more than one ion. In a Paul trap, multiple ions will push each other away from the potential minimum due to Coulomb repulsion. Away from the center, the ions will experience enhanced micromotion and, due to collisions with the other ions in the trap, RF heating which greatly increases their cooling limit. On the other hand, multiple ions at the zero-point would improve the signal to noise ratio by the square root of the number of ions

so there was interest in experimenting with other trap geometries which might make this possible. One such geometry is the linear Paul trap consisting of four parallel rods. The cross section of a linear Paul trap looks just like a conventional Paul trap, but rather than being symmetric about the axis of the endcaps, it is symmetric along the axis of the rods. Such a trap has a line along its center where the potential is zero instead of a single point. Consequently it is possible to have many ions along this line all in the zero-point energy level. Linear traps are recognized as having very good potential for frequency standards and have been investigated by a group at the Jet Propulsion Laboratory (JPL) [Pre89] and the NBS group (later known as the National Institute of Standards and Technology or NIST) among others.

In 1990 the Seattle group demonstrated the Paul-Straubel trap [Str55, Yu90] with single Ba^+ ions. This is a variant of the Paul trap in which the endcaps have been removed to infinity. There are many variations on this idea. Some Paul-Straubel traps consist of the ring of a Paul trap with the endcaps removed. Others consist of a flat strip of metal with a hole drilled through it to take the place of the ring. The Paul-Straubel trap does not have a perfect quadrupole potential but very closely approximates it at the center of the potential well. Since this is exactly where single ions reside when they are cooled, Paul-Straubel traps work quite well. While they do have increased anharmonicity and require higher voltages to get the same trap well depth, the Paul-Straubel trap has a much more open geometry. This has two advantages. First, it is easier to focus laser light into the trap with less background because there are fewer electrodes to scatter stray light. Second and most important, the simpler geometry makes smaller traps possible. A smaller trap means a higher secular frequency which can make sideband cooling and the Lamb-Dicke regime accessible where it might otherwise not have been. A third characteristic of Paul-Straubel traps which turned out to be particularly useful was that they are much easier to heat. During the ion loading process, impurities may be deposited on the trapping electrodes. These impurities can cause contact potentials which can lead to increased micromotion of the ions in the trap. Heating of the ring electrode can drive off such impurities and can be accomplished simply by running a current through the ring. On the other hand, heating of the endcaps (in a Paul trap) can't be accomplished with a current because they are only electrically contacted on one end (an electron beam is one possible alternative). In contrast, the ring electrode is often electrically

contacted on both sides and if it is thin enough a current run through it provides the desired heating. If the endcaps are no longer near the ring (as is the case with the Paul-Straubel trap) then simply heating the ring suffices. A numerical analysis of the Paul-Straubel trap for certain limiting cases was given by Brewer et. al. in 1992 [Bre92] and for more general cases by Yu et. al. in 1995 [Yu95]. Some calculations for the Paul-Straubel trap were performed by Schrama et. al. in 1993 [Sch93] who also demonstrated an endcap-only trap—that is a trap with the ring removed and the endcaps providing the approximate potential.

Yet another type of trap is the so called, “racetrack trap”. In a linear Paul trap, a weak potential must be applied to the ends to prevent ions from escaping. This requirement is relaxed in a racetrack trap where the ends of a linear trap are connected to make a loop. In either of these two types of traps (and in some other types as well), if large numbers of ions are loaded one expects to see chaotic behavior until they are cooled. As they are cooled it is conceivable that the ions may go through a phase transition to a more ordered state. In fact this was observed in 1992 by Wineland’s group [Rai92] with Hg^+ ions in a linear trap and by Walther’s group [Wak92] with Mg^+ ions in a racetrack trap (also known as a storage ring trap).

As far as high resolution spectroscopy is concerned, cooling the ion(s) into the Lamb-Dicke regime is only part of the problem. The other part is having lasers which are spectrally narrow enough to measure the width of the narrow transitions now detectable with electron shelving. Whereas the narrow transitions have on the order of 1 Hz widths, commercially available stabilization of dye lasers¹ is only to 1 MHz at best. Using the well known Pound-Drever stabilization scheme [Dre81, Dre83] it is possible to lock the laser to an external reference cavity and lock the cavity to a molecular line (or to the ion itself) for both long and short term stability in the tens of kilohertz region. By using additional thermal and vibrational isolation techniques for the cavity this can be brought down to the order of several kilohertz [Nag90]. To get long term stability below 1 kHz however seems to involve a significant increase in the amount of effort required. In 1992 Bergquist at NIST achieved *long term* laser stability of 30 Hz and was able to resolve sub 80 Hz features in single Hg^+

¹ As of this writing the only type of laser providing the desired tunability and power at the required wavelengths is the dye laser. This is changing rapidly as diode lasers generate more power and become available in wider wavelength ranges.

ions [Ber92].

A great deal of the work that has been done with trapped single ions has its roots in the series of abstracts that Dehmelt published on the Tl^+ mono-ion oscillator between 1973 and 1975 [Deh73, Deh75, Win75]. However, until recently, no one had attempted to work on Tl^+ itself or even other group IIIA ions. As has been mentioned, the advantage of group IIIA ions is the weak dependence on external fields of the clock transition which is between two states with $J = 0$. Such a transition can usually be found in non-group IIIA ions by forcing the transition to go from $m_F = 0$ to $m_F = 0$, but this requires optical pumping. Such states usually still have a quadrupole moment as well which will interact with a field gradient. States with $J = 0$ have no quadrupole moment. The difficulty with group IIIA ions is that the strongly allowed $^1S_0 - ^1P_1$ transition that would normally be the best candidate for a cooling transition is in the vacuum ultraviolet (VUV) making it difficult to synthesize and work with. Instead it is necessary to use the intercombination line, $^1S_0 - ^3P_1$. This transition is accessible (though still well into the UV), but is weaker than the strongly allowed transition and so not as efficient at cooling. This difficulty prevented much work from being done on group IIIA ions until recently when both Walther's group and Dehmelt's group began work on In^+ . In 1994 both groups successfully trapped single laser cooled In^+ ions [Pei94, Bur95]. Our work in this area is the subject of this thesis.

Chapter 2

THEORETICAL ASPECTS OF ION TRAPPING AND COOLING

In this chapter I will describe in detail some of the parameters that are important to ion trapping and cooling. When dealing with RF traps, the RF drive frequency and voltage along with the trap size together determine the pseudopotential well depth and secular frequency. The equations used to determine these values, along with a description of how they change when one goes from a Paul trap to a Paul-Straubel trap, will be given. The understanding of various heating and cooling effects for both the buffer cooled ion cloud case and the laser cooled single ion case is also essential. These topics lead to a description of zero-point confinement and Lamb-Dicke confinement, which are important for Doppler-free high resolution spectroscopy. Finally, I will describe the effect of ambient fields with respect to optical pumping and the Hanle effect.

2.1 *Trap Characteristics*

A charged particle subject to a homogeneous oscillating field, such as that between the plates of an ideal parallel plate capacitor, will feel a time dependent force but no time averaged force. On the other hand if the field is inhomogeneous, such as that in a capacitor with curved plates, then there will be a time averaged force. To see how this can be possible, consider more closely the curved plate capacitor. To be specific, let's say that the two plates have radii of curvature r_1 and r_2 and a common axis. Let there be a sinusoidally time varying potential applied across the plates. If we position the charged particle halfway between r_1 and r_2 it will experience a force whose direction and amplitude will vary at the same sinusoidal frequency and it will oscillate along a radial line between the two plates. During the half cycle where the particle is closer to the outer plate it is in a region of weaker field than during the half cycle where it is closer to the inner plate. Consequently, the average force in the outer half cycle is smaller than the average force in the inner half cycle. Since the

two forces are oppositely directed towards the center of motion, the inner half cycle force wins and the particle feels an overall time averaged force outwards towards the region of weaker field. This is the underlying principal of the RF quadrupole, or Paul trap. An RF potential is applied between the ring and endcaps which gives an RF quadrupole potential. The field amplitude is smallest in the center so a charged particle inside the region enclosed by the trap electrodes will feel an effective force pushing it towards the geometric center of the trap. In this way a charged particle can become trapped.

Following Dehmelt [Deh67] this can be made quantitative. The z-direction equation of motion is

$$m\ddot{z} = eE_0\cos\Omega t, \quad (2.1)$$

where m is the mass of the ion, e is the electron charge, E_0 is the amplitude of the field associated with the RF drive, and Ω is the drive frequency. We seek a solution to this equation of the form

$$z(t) = \bar{z} + \delta(t), \quad (2.2)$$

where \bar{z} is a constant and $\delta(t)$ is of the form $\delta_0\cos\Omega t$ and is a small deviation in position from \bar{z} . Substituting this expression for $z(t)$ into eq. 2.1 gives

$$-\delta_0\Omega^2\cos\Omega t = eE_0\cos\Omega t, \quad (2.3)$$

from which we get that

$$\delta_0 = -\frac{eE_0}{m\Omega^2}. \quad (2.4)$$

If we now assume that the field is inhomogeneous and expand it around \bar{z} to first order we get

$$E_0(z(t)) \simeq E_0(\bar{z}) + \frac{\partial E_0(\bar{z})}{\partial \bar{z}}\delta(t)\cos\Omega t. \quad (2.5)$$

The time-averaged force on the ion is then

$$\langle F(\bar{z}) \rangle = \langle eE_0(z(t))\cos\Omega t \rangle \simeq \left\langle e\frac{\partial E_0(\bar{z})}{\partial \bar{z}}\delta(t)\cos\Omega t \right\rangle. \quad (2.6)$$

Since $\langle \delta(t)\cos\Omega t \rangle = \frac{1}{2}\delta_0$ we have

$$(1/e)\langle F(\bar{z}) \rangle = -\frac{\partial E_0(\bar{z})}{\partial \bar{z}}\frac{eE_0(\bar{z})}{2m\Omega^2}. \quad (2.7)$$

This implies a pseudopotential of

$$\psi(\bar{z}) = eE_0^2(\bar{z})/4m\Omega^2 \quad (2.8)$$

($|\mathbf{F}| = -|\nabla\psi|$). This result can be extended to 3 dimensions. We can write the quadrupole potential between the trap electrodes as,

$$\phi = A_0 \cos\Omega t (x^2 + y^2 - 2z^2). \quad (2.9)$$

Taking the gradient of ϕ we can get an expression for the amplitude of the electric field in the trapping region:

$$E_0^2 = 4A_0^2(\bar{x}^2 + \bar{y}^2 + 4\bar{z}^2). \quad (2.10)$$

Substituting eq. 2.10 into eq. 2.8 gives

$$\psi(\bar{x}, \bar{y}, \bar{z}) = \frac{eA_0^2}{m\Omega^2}(\bar{x}^2 + \bar{y}^2 + 4\bar{z}^2). \quad (2.11)$$

We would like to express this in terms a well depth, \bar{D} :

$$\psi(\bar{x}, \bar{y}, \bar{z}) = \bar{D}[\bar{x}^2 + \bar{y}^2 + 4\bar{z}^2]/(4z_0^2). \quad (2.12)$$

Equating eq. 2.12 to eq. 2.11, we can solve for \bar{D} :

$$\bar{D} = \frac{eA_0^2 4z_0^2}{m\Omega^2}. \quad (2.13)$$

The quantity \bar{D} is referred to as the pseudopotential well depth. When designing a quadrupole trap, it is desirable to have the well depth in the z direction be the same as the well depth in the r direction. This prevents an ion, that may be adequately trapped in one direction, from hitting an electrode in the other direction. By examining eq. 2.12, we see that this requirement is satisfied if $r_0 = 2z_0$. Note that this is different than the similar condition for a Penning trap where $r_0 = \sqrt{2}z_0$. This is because in the Penning trap potential, the coefficient of z is 2, whereas in the Paul trap *pseudopotential*, the coefficient of z is 4. Substituting $r_0 = 2z_0$ into eq. 2.13 gives

$$\bar{D} = \frac{eA_0^2 r_0^2}{m\Omega^2}. \quad (2.14)$$

From the expression for \bar{D} we can now derive an expression for the trap axial secular frequency, ω_z , by noting that

$$e\bar{D} = \frac{mz_0^2\omega_z^2}{2}. \quad (2.15)$$

By using eq. 2.8, eq. 2.10, and the fact that $\bar{D} = \psi(\bar{z} = z_0)$ we have

$$e\bar{D} = \frac{e^2 E_0^2(\bar{x} = 0, \bar{y} = 0, \bar{z} = z_0)}{4m\Omega^2} = \frac{4e^2 A_0^2 z_0^2}{m\Omega^2}. \quad (2.16)$$

Putting together equations 2.15, and 2.16 we have

$$\omega_z^2 = 8 \frac{e^2 A_0^2}{m^2 \Omega^2}. \quad (2.17)$$

From eq. 2.9 the amplitude of the RF drive voltage is

$$V_0 = \phi(r = 0, z = z_0, t = 2\pi/\Omega) = -2z_0^2 A_0. \quad (2.18)$$

Substituting this into eq. 2.17 we can express the secular frequency in terms of the voltage:

$$\omega_z^2 = 2 \frac{e^2 V_0^2}{m^2 \Omega^2 z_0^4}. \quad (2.19)$$

Finally using the assumption that $r_0 = 2z_0$ we have

$$\omega_z = 4\sqrt{2} \frac{eV_0}{m\Omega r_0^2}. \quad (2.20)$$

2.2 Optimal Trap Geometry

In the geometry of a classical Paul trap, it is usually the case that $r_0 \approx z_0$; that is, the endcap separation is about the same as the ring diameter. For a Paul-Straubel trap this condition can't hold because the endcaps have been effectively removed to infinity. As a result this geometry only approximates a quadrupole potential very near the center of the trap. Away from the center the anharmonic terms in the potential become very significant and lead to amplitude and position dependent secular frequency shifts. Such anharmonicities exist in Paul traps as well, at least due to fabrication imperfections, but most likely because the trap electrode cross

sections are usually not hyperbolic. In a Paul-Straubel trap, more voltage is generally required to achieve the same well depth—this is called the voltage loss factor. Various geometries lead to different anharmonicities and voltage loss factors. Usually there is a trade off between them; they can't both be optimized at the same time. For single ions which are laser cooled and reside very close to the center of the potential, anharmonicity is probably less important than voltage loss. However the ion doesn't start off laser cooled. When it is first loaded, it may be quite hot and anharmonicities may interfere with initial cooling. Consequently, a compromise between these two competing effects is used.

Yu's numerical analysis of trap geometries [Yu95] shows that for a toroidal Paul-Straubel trap an ideal loss factor of 3.6 is reached for a trap aspect ratio, $A = t/r_0 = .17$ where t is the thickness of the ring and r_0 is the inner diameter. The anharmonic terms for this aspect ratio are quite large. The aspect ratio corresponding to the optimal value for C_4 (the fourth order term) is $A = 1.1$ (loss = 8.7!) and that for the optimal value of C_6 (the sixth order term) is $A = .75$ (loss = 5.6). See Section 4.9 for details of the trap geometry used in this experiment.

2.3 *Trapping Ion Clouds*

In the initial stages of the experiment, large ion clouds were trapped. This was done as a proof of principle before moving on to the more difficult single ion work. It was shown in 1978 [Neu78] that small ion clouds (< 50) could be laser cooled, however for our initial work on In^+ we did not want to depend on laser cooling. Using a buffer gas to cool instead, it is possible to trap and cool to room temperature, on the order of 10^5 ions leading to a larger signal.

2.3.1 *Buffer Gas Cooling*

Buffer gas cooling works by diffusing a gas into the trapping region which is at a substantially lower temperature (usually room temperature) than the ion cloud. Collisions between the ions and the buffer gas tend to thermalize the ions with the gas. Thus the lowest temperature of the cloud is limited by the buffer gas temperature. An estimate of the cooling rate is derived by analyzing head-on collisions between two particles of different mass and temperature. Let $r = m/M$ be the mass ratio

where m is the mass of the buffer gas and M is the mass of the ion. The conservation equations give:

$$\frac{1}{2}mv_0^2 + \frac{1}{2}MV_0^2 = \frac{1}{2}mv_1^2 + \frac{1}{2}MV_1^2 \quad (2.21)$$

and,

$$mv_0 + MV_0 = mv_1 + MV_1 \quad (2.22)$$

where the subscripts 0 and 1 refer to initial and final respectively. To first order in the value of $\sqrt{1-2r}$ these equations can be solved to give:

$$v_1 = (2V_0 - v_0)(1 - r) \quad (2.23)$$

$$V_1 = V_0 - 2(V_0 - v_0)r. \quad (2.24)$$

If we take the change in the energy of the ion due to a collision with a buffer gas particle to be:

$$\delta E = \frac{1}{2}MV_1^2 - \frac{1}{2}MV_0^2 \quad (2.25)$$

then we can also write,

$$\delta E = -2\left(\frac{v_0}{V_0} + 1\right)MV_0^2r. \quad (2.26)$$

If the ions are initially quite hot then we may take $V_0 \gg v_0$ and,

$$\delta E \simeq -4\left(\frac{1}{2}MV_0^2\right)r \quad (2.27)$$

or,

$$\frac{\delta E}{E} \simeq -4\frac{m}{M}. \quad (2.28)$$

In the case of In^+ ions cooled by a He buffer gas the mass ratio is 0.03486 which gives:

$$\frac{\delta E}{E} \simeq 0.140 \quad (2.29)$$

as the fractional energy loss per collision.

Next we want to calculate the collision rate between each In^+ ion and the ambient He buffer gas. In the reference frame of an In^+ ion the flux of He passing by is $n_H\bar{V}$, where n_H is the number density of He and \bar{V} is the mean relative velocity of the In^+ ions and the He atoms. The fraction of this flux that is scattered is given by $n_H\bar{V}\sigma_0$ where σ_0 is the scattering cross section between In^+ ions and He atoms. Since this

is the number of He atoms scattered *per second*, it is the desired rate. Note that the rate that we are interested in is independent of the density of In^+ ions. If there are more In^+ ions then there are more *total* collisions, but there is also more *total* energy to carry away by the same proportion. The He density is given simply by:

$$n_H = \frac{P_H}{kT_H} \quad (2.30)$$

where P_H is the pressure of He and T_H is its temperature. The mean relative velocity is given approximately by adding the separate mean velocities in quadrature:

$$\bar{V} \approx \sqrt{\bar{v}_I^2 + \bar{v}_H^2}. \quad (2.31)$$

These separate mean velocities can be given in terms of the respective temperatures [Rei65, p. 268]:

$$\bar{v}_i = \sqrt{\frac{8kT_i}{\pi m_i}}. \quad (2.32)$$

The cross section of two hard spheres is given by $\sigma_0 = \pi(r_1 + r_2)^2$. For In^+ and He we can take these two radii to be approximately equal so that,

$$\sigma_0 \approx 4\pi r^2. \quad (2.33)$$

Putting all of this together we get a collision rate of:

$$R = \frac{4\pi r^2 P_H}{kT_H} \sqrt{\frac{8k}{\pi} \left(\frac{T_I}{M} + \frac{T_H}{m} \right)}. \quad (2.34)$$

Taking $T_H = 300$ K, $T_I = 3100$ K (roughly 1/10 the well depth), $M = 114.82$ amu, $m = 4.003$ amu, $r = 0.53 \times 10^{-10}$ m, and $P_H = 10^{-4}$ Torr we get $R = 1.66$ s⁻¹. Taking this together with the fraction of energy carried off per collision (for hot In^+ , use eq. 2.29) we have that the In^+ ions will be brought into thermal equilibrium with the He buffer gas in a time of order 4 s. The actual time to thermal equilibrium may be substantially faster since, as can be seen from eq. 2.25, buffer gas cooling gets more efficient as the ion is cooled.

2.3.2 Heating Effects

There are several possible heating effects. One possibility is collisions with the trap electrodes, but this effect is self limiting since only the hottest ions make such collisions and once they do so they are usually ejected from the trap. The primary heating

mechanism is called RF heating. If one considers just two ions in a RF trap then the equation of motion for each ion is:

$$m\ddot{x}_1 = -f(x_1)\cos\Omega t + \frac{kq^2}{(\bar{x}_1 - \bar{x}_2)^2} \quad (2.35)$$

$$m\ddot{x}_2 = -f(x_2)\cos\Omega t - \frac{kq^2}{(\bar{x}_2 - \bar{x}_1)^2} \quad (2.36)$$

where x_i is the position of each ion, f is the position dependent force of the RF drive at frequency Ω and the second term in each equation is the Coulomb interaction. The difference of these two equations represents internal motion of the two ions:

$$m(\ddot{x}_1 - \ddot{x}_2) = -(f(x_1) - f(x_2))\cos\Omega t + \frac{2kq^2}{(\bar{x}_1 - \bar{x}_2)^2}. \quad (2.37)$$

Since the field strength is inhomogeneous and the Coulomb collisions are at random times and in random directions, the first term is a randomly changing force which couples secular motion to internal ion motion. Essentially the ions are able to draw energy from the RF drive, hence the name RF heating. Although I have not calculated the size of this heating effect, experience shows that it is small compared to buffer gas cooling for ion clouds as is evidenced by how close buffer gas cooling brings the ion cloud to room temperature. RF Heating is virtually non-existent for single ions since there are no collisions (except with residual gas atoms whose effect is closer to that of buffer gas heating, than to RF heating).

2.3.3 Number of Ions Trapped

The maximum number of ions in a Paul trap can be calculated from the volume marked off by the electrodes taken together with the charge density obtained by requiring that the energy of Coulomb repulsion be equal to the well depth of the trap [Deh67]. In the Ba^+ experiment, buffer gas cooled clouds were observed to occupy a space equal to about 1/3 of the endcap separation. This corresponds to pseudopotential energy of about 1/10 the well depth. Using the same Coulomb energy argument, but now equating this to $e\bar{D}/10$ we get a maximum Paul trap capacity of:

$$N_{max} = 2.57 \times 10^5 \bar{D} z_0 \quad (2.38)$$

where \bar{D} is in volts and z_0 (1/2 the endcap separation) is in cm. For example, our Paul trap had $\bar{D} = 26$ eV and $z_0 = 0.0805$ cm which gives $N_{max} \approx 540,000$ ions. Estimates based on observed fluorescence are actually about an order of magnitude less, so this number is to be taken as very approximate. Still, eq. 2.38 can be used effectively to determine how the capacity *varies* with \bar{D} and z_0 , if not to determine the absolute capacity itself.

2.4 Trapping Single Ions

For trapping of single ions, the various heating effects become weak enough that laser cooling works. A buffer gas would not be effective in cooling single ions because the temperature would get no lower than 300 K which corresponds to a Doppler width that would make the ion fluorescence unobservable. To observe the ion's fluorescence, the Doppler width must be on the order of the natural width.

2.4.1 Laser Cooling

While buffer gas cooling provides an efficient way to cool hot ions down to room temperature, laser cooling provides a very efficient way of cooling ions into the sub-mK realm. Here I will give the various cooling rates and limits appropriate to our experiment. There are two cooling regimes leading to different limits: the weak binding regime where $\gamma \gg \omega_z$ ¹ and the strong binding regime where $\gamma \ll \omega_z$.

In the weak binding regime where the atoms are approximately free, Wineland and Dehmelt [Win75, Win79a] calculate the minimum energy to be:

$$\langle E_{min} \rangle_{wb} = \frac{\hbar\gamma}{4} \quad (2.39)$$

where γ is the natural width of the transition. If we were to have an ensemble of atoms with this energy that were weakly interacting so that they were thermalized, the equivalent minimum temperature would be:

$$T_{min}^{wb} = \frac{\hbar\gamma}{2k}. \quad (2.40)$$

¹ I will refer to the z-component for all values where it matters, though it should be understood that I am simply leaving out discussion of the other components for simplicity.

This is also known as the Doppler cooling limit. The minimum temperature is limited by the recoil energy $R = (\hbar k)^2/2M$.

For the trapping parameters that we use, In^+ is in the strong binding regime ($\gamma \ll \omega_z$) because $\gamma = 2\pi \times 360kHz$ and $\omega_z = 2\pi \times 1MHz$. For strong binding, the ion fluorescence takes on resolvable sidebands at the secular motion frequency. The laser is tuned to the first lower sideband for cooling, hence this is also called sideband cooling. The field seen by the atom is:

$$E_a = E_0 \sin(kx - \omega t) \quad (2.41)$$

where $x = x_a \sin(\omega_z t)$ and x_a is the amplitude of the ion's motion. The field becomes:

$$E_a = E_0 \sin(kx_a \sin \omega_z t - \omega t) \quad (2.42)$$

which corresponds to frequency modulation at ω_z and with modulation index, $\beta = kx_a$. This can be expanded in Bessel functions to give the sidebands. The intensity of the m^{th} sideband is proportional to $J_m^2(kx_a)$. If the laser is tuned to the m^{th} lower sideband, light of energy $\hbar(\omega_0 - m\omega_z)$ is absorbed and light at $\hbar\omega_0$ is emitted leading to an overall cooling process. The cooling rate given by Wineland is:

$$\frac{dE}{dt} = \hbar m \omega_z \frac{I \sigma_0}{\hbar \omega_0} J_m^2(kx_a) \quad (2.43)$$

where m is the order of the Bessel function that is largest for the given modulation index, I is the laser intensity and $\sigma_0 = \lambda^2/2\pi$ is the light scattering cross section (here we are assuming an allowed transition).

What sort of cooling rate do we have for an ion which is initially at room temperature? This temperature corresponds to a secular motion amplitude of $x_a = \sqrt{k_B T / M \omega_z^2} = 2.34 \times 10^{-3}$ cm. For a cooling laser at 231 nm this gives a modulation index of $\beta = kx_a = 636$. Maximum cooling will be achieved by tuning the laser to the sideband with the largest amplitude. Even if the sidebands aren't resolved at this point, this will tell us the appropriate detuning frequency. For this modulation index the $m = -629$ (the negative sign indicates red detuning) order Bessel function is largest which suggests a detuning of 629 MHz (it should come as no surprise that this is roughly half of the Doppler width). We have $J_{629}^2(636) = 6.18 \times 10^{-3}$. The quantity $I \sigma_0 / \hbar \omega_0$ corresponds to the transition rate which for In^+ begins to saturate

at about 1 MHz, so assume that the laser intensity I is such that the rate is 1 MHz (less than $50 \mu\text{W}$ of UV is required to do this). Putting all of this together gives $dE/dt = -1.61 \times 10^{-2} \text{ eV/s}$, or almost -400 K/s . So one would expect the ion to reach its cooling limit within one second. Typically the cooling process involves loading the ion and then “manually chirping” the laser frequency up from its initial detuning as the ion cools (as the ion cools the optimal detuning decreases). For this situation it is clear that the “chirp” can be quite fast by human standards.

The assumption that the ion is initially at room temperature after loading is completely arbitrary. In fact it could be a good deal hotter because the initial temperature of the ion is determined by where it is in the pseudopotential at the moment that it is ionized and how fast it is moving at that point. Since the ion can be created on the edge of the trapping potential, it can have an initial temperature comparable to the well depth. The speed that the atom is moving at when it is ionized is determined by the thermal distribution of atoms being emitted from the oven. The well depth in our Paul trap was 26 eV whereas the oven is about $1,000^\circ\text{C}$ corresponding to a kinetic energy of only a fraction of an eV. Therefore the initial position of the ion in the pseudopotential should be the dominating factor in determining initial temperature. In either case, this shows that the initial temperature can be much greater than room temperature. Of course only atoms with less energy than the trap well depth when they are ionized in the trapping region, will remain trapped. If we carry through the same analysis with ions that start out with an energy corresponding to $e\bar{D}/10$ (1/10 the well depth) we find that $\beta = 6,364$, $J_{6350}(\beta) = 0.0363$. This suggests a detuning of 6.3 GHz. If we still assume this modulation index but only detune the laser by 1 GHz then $J_m(\beta) = .01$ and $dE/dt = -2.62 \times 10^{-3} \text{ eV/s}$ or about -32 K/s ! Cooling would be even slower for less detuning. If this were the initial detuning one would want to leave the laser at this frequency for some time (about 190 s in this case) before scanning the laser frequency up. Ultimately we did find that a precooling period of about 40 s was necessary before scanning the laser. As has just been shown this is easily explained by assuming that the initial ion temperature is somewhat above 300 K.

A quantum mechanical treatment of sideband cooling [Win79a] in the strong binding regime which takes into account the effect of recoil on the emission and absorption

spectra, gives a cooling rate of:

$$\langle dE/dt \rangle^{SB} = \frac{I\sigma_0}{\hbar\omega_0} R(5/16(\gamma/\omega_z)^2 - \langle n_z \rangle) \quad (2.44)$$

where $\langle n_z \rangle$ is now the mean occupation number of the quantized potential well and R is the recoil energy. We find the cooling limit by setting this equal to zero and get:

$$\langle n_z \rangle_{min}^{SB} = 5/16(\gamma/\omega_z)^2. \quad (2.45)$$

For In^+ we have $\langle n_z \rangle_{min}^{SB} = 0.04$ which is virtually in the zero point energy level. For this reason it is thought that In^+ will cool directly to zero point confinement. By way of comparison, the Ba^+ cooling transition has a natural width of about 15 MHz and unless the secular frequency is dramatically increased (a non-trivial task), the mean occupation number, $\langle n_z \rangle$, will be substantially greater than 1.

2.4.2 Optimal Detuning for Laser Cooling

As discussed in the previous section, the optimal initial laser detuning is approximately one half of the Doppler width or on the sideband with largest amplitude. However since the initial temperature and Doppler width of a single ion can only be guessed at, some precooling (leaving the laser at the detuned frequency without scanning up) may be required before the Doppler width reaches the guessed value. A room temperature In^+ ion will be cooled most efficiently when the initial laser detuning is set to about 630 MHz.

2.4.3 Lamb-Dicke Confinement

Zero-point confinement refers to a restriction of the energy level of the ion to the zero point energy. On the other hand, Lamb-Dicke confinement [Dic53] refers to the restriction of ion secular motion to a region smaller than the wavelength of the radiation being used to excite it divided by 2π . Previously we saw that the modulation index due to the ion's motion was $\beta = 2\pi x_a/\lambda$. The amplitude of the first sideband is proportional to $\beta/2$ so its intensity is proportional to $\beta^2/4$ or $(\pi x_a/\lambda)^2$. Therefore, if the ion is in the Lamb-Dicke regime, then $x_a \ll \lambda$ and the sidebands are small compared to the carrier. Thus, the importance of Lamb-Dicke confinement is the

complete elimination of the first order Doppler shift. In theory the natural lineshape is completely revealed in the carrier.

If we imagine a weakly coupled thermalized ensemble of In^+ ions cooled to the sideband limit then $\langle n_z \rangle h\nu_z = \frac{1}{2}kT$ and using $\langle n_z \rangle = \langle n_z \rangle_{min}^{SB}$ we have $T_{min} = 3.8\mu K$. We can set this temperature equal to $M\omega_z^2 x_a^2/k$ to get the secular motion amplitude as a function of mean occupation number:

$$x_a = \sqrt{\frac{2\langle n_z \rangle h}{M(2\pi)^2 \nu_z}}. \quad (2.46)$$

Substituting values we have $x_a = 2.6$ nm, well below the wavelength of 231 nm and well into the Lamb-Dicke regime.

2.4.4 Heating Effects

For a single ion in an RF trap maintained at ultra-high vacuum there are very few heating effects. This is evidenced experimentally by the fact that single ions have been maintained in traps without laser cooling for 12 hours or more and that this has been repeated *on the same ion* for weeks at a time. Since there are no other ions in the trap there is no RF heating to speak of and since the buffer gas pressure has been reduced by 6 orders of magnitude or more, buffer gas heating is all but non-existent. If more than one ion is in the trap however, RF heating does become a factor. Experience with other ions has shown that with as few as 2 or 3 ions in a Paul-Straubel trap, the lifetime may be as short as several minutes.

2.5 Expected Fluorescence Signal Strength

It is useful to have an estimate on the expected fluorescence signal strength. For laser power incident on an ion cloud well below saturation, the emission rate is

$$R = B\rho(\omega) \quad (2.47)$$

where B is the Einstein B coefficient and $\rho(\omega)$ is the spectral energy density of the laser at frequency ω . B can be determined from the lifetime of the transition:

$$B = \frac{\pi^2 c^3}{\hbar \omega_0^3} A = \frac{\pi^2 c^3}{\hbar \omega_0^3} 2\gamma \quad (2.48)$$

where $A = 2\gamma = 1/\tau$ (τ is the atomic lifetime and γ is its natural width). The spectral energy density is [Cor77, p. 284]

$$\rho(\omega) = \frac{4\pi}{c} I(\omega) \quad (2.49)$$

where we can approximate the Lorentzian laser lineshape with a Gaussian:

$$I(\omega) = I_0 \frac{2}{\Delta_L \pi^{1/2}} e^{-\frac{4(\omega-\omega_0)^2}{\Delta_L^2}}. \quad (2.50)$$

The Doppler broadening is the Gaussian:

$$D(\omega) = \frac{2}{\Delta_D \pi^{1/2}} e^{-\frac{4(\omega-\omega_0)^2}{\Delta_D^2}}. \quad (2.51)$$

If we now also approximate the natural lineshape by a Gaussian as well, then the width of the convolution is simply the sum of the individual widths in quadrature. Let the natural lineshape be

$$N(\omega) = \frac{2}{\Delta_N \pi^{1/2}} e^{-\frac{4(\omega-\omega_0)^2}{\Delta_N^2}}. \quad (2.52)$$

The convolution is now

$$C(\omega) = I_0 \frac{2}{\Delta_c \pi^{1/2}} e^{-\frac{4(\omega-\omega_0)^2}{\Delta_c^2}} \quad (2.53)$$

where $\Delta_c^2 = \Delta_N^2 + \Delta_D^2 + \Delta_L^2$ (note that each Δ_i is a 1/e width; the FWHM is found from $\Delta\omega_i = \Delta_i \sqrt{\ln 2}$). Putting this all together, we get the fluorescence rate

$$R(\omega) = \frac{4\pi}{c} I_0 \frac{\pi^2 c^3}{\hbar \omega_0^3} 2\gamma \frac{2}{\Delta_c \pi^{1/2}} e^{-\frac{4(\omega-\omega_0)^2}{\Delta_c^2}} \quad (2.54)$$

The maximum signal is taken when $\omega = \omega_0$:

$$R_{max} = \frac{4\pi}{c} I_0 \frac{4\pi^2 c^3}{\hbar \omega_0^3 \pi^{1/2}} \frac{\gamma}{\Delta_c}. \quad (2.55)$$

After converting to the half maximum widths and performing some algebra, we have

$$R_{max} = I_0 \frac{\sqrt{\ln 2} \lambda^3}{\pi^{1/2} \hbar c} \frac{\gamma}{\sqrt{\gamma^2 + \Delta\omega_D^2 + \Delta\omega_L^2}} \quad (2.56)$$

where $\Delta\omega_D$ is the Doppler width and $\Delta\omega_L$ is the laser width. This rate is expressed in counts per second per ion irradiated. The collection efficiency is considered in Section 4.10.2 and an estimate of the detected signal is determined in Section 3.2.2.

2.6 Saturation Power and Power Broadening

Even though the previous calculation was done assuming that the laser power was well below saturation, we can get an idea of where the saturation power is by letting

$$R_{max} = 2\gamma \quad (2.57)$$

and solving for I_0 . This gives

$$I_0^{sat} = \frac{\pi^{1/2} \hbar c \sqrt{\gamma^2 + \Delta\omega_D^2 + \Delta\omega_L^2}}{\sqrt{\ln 2} \lambda^3}. \quad (2.58)$$

To get the saturation power we simply multiply by the laser spot area: $\pi\omega_0^2$ where ω_0 now refers to the spot size.

$$P_{sat} = \frac{\pi^{3/2} \hbar c \omega_0^2 \sqrt{\gamma^2 + \Delta\omega_D^2 + \Delta\omega_L^2}}{\sqrt{\ln 2} \lambda^3}. \quad (2.59)$$

In addition to the previously considered widths there is also that due to power broadening. Following Loudon [Lou83, p. 65, eq. 2.123], the power broadening is:

$$\Delta\nu_{PB} = 2(\gamma^2 + \frac{1}{2}|V|^2)^{1/2} \quad (2.60)$$

where $2\gamma = 1/\tau$ is the natural width ², and $|V| = eE_0|D_{12}|/\hbar$ is the matrix element of the transition. We can rewrite the matrix element in terms of the lifetime as:

$$|D_{12}| = \sqrt{\frac{1}{\tau} \frac{3\pi\epsilon_0 \hbar c^3}{e^2 \omega_0^3}}. \quad (2.61)$$

Since laser intensity $I = \langle S \rangle = c\epsilon_0 E_0^2/2 = P/A$, we have $E_0^2 = 2P/Ac\epsilon_0$. Putting this all together gives,

$$\Delta\nu_{PB} = 2\gamma \sqrt{1 + \frac{3\lambda^3 \tau P}{2\pi^2 c \hbar A}}. \quad (2.62)$$

As an example, if $P = 50 \mu\text{W}$ and $A = \pi(40\mu\text{m})^2$ then $\Delta\nu_{PB} = 14.8 \text{ MHz}$, almost 50 times the natural width!

² My apologies to the reader for the inconsistent definitions of γ at different locations in the text. In all cases it is explicitly stated which definition is being used.

2.7 Quantum Jumps

Quantum jumps, or electron shelving, is the technique which makes it possible to see the highly infrequent transitions from a metastable state by the presence or absence of fluorescence from a strong transition sharing the same ground state. Typically, one laser is used for each transition, however, two lasers are not required. In the case of In^+ the metastable state has a lower energy than the excited state of the strong transition. While repeatedly exciting the ion into this excited state, it usually decays into the ground state. Occasionally however, it will decay into the metastable state. When this happens the fluorescence from the strong transition will shut off, just as in the normal electron shelving scheme. Such shelving cannot be used to do spectroscopy on the narrow transition because the number of times the electron is shelved per unit time has no relationship to laser frequency. It can be used to demonstrate the presence of a single ion and measure the lifetime of the metastable state. This is useful in the initial stages of the experiment before a clock laser is set up.

In order to observe quantum jumps with just one laser, there are certain requirements on the amount of fluorescence being emitted by the strong transition. First, the detection integration time per channel must be shorter than the average lifetime of the weak transition. If it isn't, then fluorescence will always be getting into the channel where the shelving took place. At best a slight reduction in signal will occur which may not be larger than the noise. Let S be the number of fluorescence counts detected in 1 second, B be the instrumental scattering background in 1 second and let us assume that we have shot noise, $N = \sqrt{B}$. If we now change our integration time per channel from 1 second to t_c , the new values for signal, background and noise will be: $S' = St_c$, $B' = Bt_c$ and $N' = \sqrt{Bt_c}$. This gives a new signal to noise ratio of $S'/N' = \sqrt{t_c}S/N$. In order to be able to detect the difference between shelved and not shelved, we need $S'/N' \geq 1$. If we set the integration time per channel to $\tau/3$ (1/3 of the weak transition lifetime) then:

$$\frac{S'}{N'} = 1 = \sqrt{\frac{\tau}{3}} \frac{S}{N}. \quad (2.63)$$

To make this concrete, we routinely observe a background of $B = 20s^{-1}$ and the lifetime of the weak transition is 140 ms. These values require $S \geq 430s^{-1}$. Anticipating the result of this thesis, that the fluorescence from a single ion detected by

the apparatus described here resulted in 100 s^{-1} of signal, we need a 4x improvement in signal to observe quantum jumps.

One way to increase signal is to stabilize the laser. Doing so increases the spectral density by the same factor that the laser width is reduced. When the laser width is larger than the width of the transition being excited, the laser power contained in the portion of the laser width that doesn't overlap the transition width is being thrown away. Thus, as the laser width is made smaller, the power required to obtain the same signal also goes down by the same factor *up until* the laser width becomes smaller than the transition width. At this point, laser power is no longer being wasted. Consequently, the same signal should be observable with the overall power reduced by a factor equal to the ratio of the unstabilized laser width to the larger of the stabilized laser width or the transition width. The reduced power means reduced background; again by the same factor. The value of $B = 20 \text{ s}^{-1}$ was observed with the laser commercially stabilized to about 5 MHz (as measured against a Tellurium saturation line). The locking of this laser to an external reference cavity and the Tellurium molecular line results in a laser linewidth of 30 kHz or a factor of over 100x improvement. This is well below the natural width of the transition, 360 kHz, so the overall factor for power reduction without signal loss is determined by the natural width and is about 10. This brings the background down to 2 s^{-1} which is comparable to the detector dark current level.

This is still not quite good enough though, because the ion is cold. Even with a 10 fold reduction in power we are still near saturation (these signal and background values were observed with about $50 \mu\text{W}$ of UV power). While background is eliminated, signal doesn't increase with the higher spectral density. With an integration time of $\tau/3$ we still only have roughly 3 counts of signal per channel which may not be observable above the dark current of the detector. This enhancement alone may still work, however, because the stable laser is able to "sit" closer to the peak of the transition without heating the ion. If this still doesn't give the extra factor required, then switching the output optic from an $f/2$ mirror to an $f/1.4$ corrected lens almost certainly will.

2.8 Ambient Magnetic Fields and Optical Pumping

The previously calculated HF Zeeman splitting was 382 kHz/Gauss in the 3P_1 state. With a commercially stabilized dye laser ($\Delta\nu_{laser} = 10$ MHz in the UV) this splitting isn't resolved. However, with a 30 kHz externally stabilized laser this splitting becomes resolved and leads to a potential optical pumping problem. If no attempt is made to shield the ambient magnetic field of roughly 1 Gauss and the incident light is linearly polarized, the ion may decay with $\Delta m_F = \pm 1$. If this happens, the transition frequency will have shifted by about 380 kHz (the ground state HF Zeeman splitting is only 1 kHz) and since this is substantially larger than the laser width, the laser is no longer resonant with the transition!

To resolve this problem, the incident light must be made circular polarized so as to excite only a $\Delta m_F = \pm 1$ transition (where the sign is determined by whether left or right circular polarized light is used). This together with using the HF line, $F = 9/2 - F = 11/2$ turns the transition into the 2-state system: $m_F = \pm 9/2 - m_F = \pm 11/2$. This conversion is only as good as the optical element used to create the circular polarized light (not to mention the purity of the input linear polarization from the laser). Any deviation could possibly excite a $\Delta m_F = 0$ transition, from $m_F = -9/2$ to $m_F = -7/2$, for instance, which would be free to decay to the $m_F = -7/2$ state, again destroying resonance. A factor of 100x reduction of ambient field down to the 10 mGauss region using Helmholtz coils would once again make all of the sublevels effectively degenerate. Alternatively, Helmholtz coils could be used to *apply* a field which would move the splitting out to a point where accidental excitations into the wrong state would be very unlikely. A 2-state system is still desirable, so both the circular polarization and the field shielding/enhancement should be used.

2.9 Hanle Effect Rotations and Signal Strength

A possible concern for fluorescence detection is the effect of angular momentum precession due to a magnetic field. Because the angular momentum of the ion precesses in whatever magnetic field happens to be present, the polarization of the emitted fluorescence may not be the same as that of the exciting laser beam. If the fluorescence were passed through a polarization analyzer, the signal strength as a function of magnetic field strength would trace out a resonance lineshape. This is usually referred

to as the Hanle effect and is often used to measure transition lifetimes. A polarizer is placed on the detector and classically, as the fluorescence polarization is rotated the detected signal is attenuated. How much rotation occurs depends on the field strength and the lifetime. The rotation takes place at the Larmor frequency so the lifetime determines how far, on average, the angular momentum will precess before emission takes place. A knowledge of the detected signal, the incident polarization, the field strength and its direction will then give the transition lifetime.

The Hanle effect can also cause a problem for detection even when a polarizer isn't present. Let the ambient magnetic field be parallel to the incident laser beam direction and call this the x direction. Let the incident polarization be in the z direction and the detection be in the y direction. That is, a line from the radiating ion to the detector is along the y axis. For this configuration, precession will be in the y-z plane. If there were no field present, there would be no precession and the detector would be ideally positioned to detect a dipole radiation pattern. With a field, however, the precession could proceed to the point where the polarization of the fluorescence is along the y axis and would not be observable along this axis. In some sense, the direction of observation is taking the place of the output polarizer (analyzer). This could easily result in a greatly attenuated signal which would be disastrous for a signal that may already be quite weak. Of course the signal will never be completely extinguished due to the finite width of the transition and averaging processes. To make this more precise I will calculate the maximum and minimum fluorescence expected in a 1 Gauss field (typical for our unshielded setups) as a percentage of the fluorescence with no field present.

For the standard Hanle effect the intensity of atomic fluorescence can be written as [Cor77, p. 483]

$$I(\alpha) \sim \frac{1}{\gamma} + \frac{\gamma \cos(2\alpha) + 2\omega_l \sin(2\alpha)}{\gamma^2 + 4\omega_l^2} \quad (2.64)$$

where α is the angle of the analyzer relative to the incident polarization, in this context $\gamma = 1/\tau$, and

$$\omega_l = g_F \frac{eB}{2m_e} \quad (2.65)$$

with g_F being the g-factor for an ion with total angular momentum F, e is the electron charge, B is the field strength and m_e is the mass of the electron. In our case we can take α to be the angle between the incident light polarization and the y-axis. From

a simple dipole radiation pattern, we would expect to see maximum signal when $\alpha = \pi/2$ and minimum signal when $\alpha = 0$. If no field is present, eq. 2.64 reduces to

$$I_0 \sim 2/\gamma. \quad (2.66)$$

For $\alpha = 0$ we have,

$$I(0) \sim \frac{1}{\gamma} + \frac{\gamma}{\gamma^2 + 4\omega_l^2} \quad (2.67)$$

where in each equation the proportionality is the same.

To calculate the appropriate Larmor frequency we need the g-factors:

$$g_F \approx g_J \frac{F(F+1) + J(J+1) - I(I+1)}{2F(F+1)} \quad (2.68)$$

and

$$g_J = \frac{J(J+1) + L(L+1) - S(S+1)}{2J(J+1)} + g_S \frac{J(J+1) - L(L+1) + S(S+1)}{2J(J+1)}. \quad (2.69)$$

For the $In^+ \ ^1S_0 - ^3P_1 \ 9/2 - 11/2$ transition, $S = 1$, $L = 1$, $J = 1$, $I = 9/2$ and $F = 11/2$. These values give $g_F = 0.273$. From eq. 2.65 $\omega_l = 2.40 \times 10^6$ ($B = 1g$). Note that this value of ω_l is very close to the decay rate, $1/\tau$, so we expect the Hanle effect to be appreciable. In general if $\omega_l \sim 1/\tau$ there is an effect and if $\omega_l \ll 1/\tau$ there is no effect. Here we have, $\gamma = 1/\tau = 2.27 \times 10^6$. Putting these values into eq. 2.66 and 2.67 with $\alpha = 0$ and then taking the ratio we get

$$\frac{I(0)}{I_0} = \frac{1}{2} + \frac{1}{2(1 + 4\omega_l^2/\gamma^2)} = 0.592. \quad (2.70)$$

That is, if the incident polarization is perpendicular to the y-axis then the observed fluorescence will be 59.2% of what it would be with no field present.

The value just obtained is somewhere in between the maximum and minimum possible intensity percentages (as a function of α). To find out what the worst case is we must calculate these maximum and minimum values. To do this we find extremum for the expression (see eq. 2.64)

$$\gamma \cos(2\alpha) + 2\omega_l \sin(2\alpha). \quad (2.71)$$

This gives a local minimum at $\alpha_{min} = 122.4^\circ$ and a local maximum at $\alpha_{max} = 32.4^\circ$. From eq. 2.64 we have for any value of α ,

$$\frac{I(\alpha)}{I_0} = \frac{1}{2} + \frac{2\cos(2\alpha) + \beta\sin(2\alpha)}{1 + 4\beta^2} \quad (2.72)$$

where $\beta = \omega_l/\gamma$. Substituting values for ω_l , γ , α and β we get:

$$\frac{I(\alpha_{max})}{I_0} = 0.714 \quad \text{and} \quad \frac{I(\alpha_{min})}{I_0} = 0.286. \quad (2.73)$$

Thus, the price for not aligning the input polarization correctly with respect to the output detection can be up to a factor of 2 in signal. Of course, this assumes a roughly 1 G field in the x direction which turned out to be the case, but certainly isn't necessarily so. In this case the best angle for the incident polarization is 32.4° from the z-axis, but if the polarization is parallel to the z-axis (the typical situation) there is only an 18% reduction in signal.

Chapter 3

INDIUM

Due to their high degree of confinement and isolation, the potential for high resolution spectroscopy on trapped single ions is quite good. The question of which ion to trap still remains. One criterion is that the ion must have transitions that are reasonably accessible to current laser technology (for optical transitions). Another criterion pertains to the strength of the proposed cooling transition and the narrowness of the proposed clock transition. A typical candidate for cooling would be a strong electric dipole transition which can supply a large amount of cooling power during the initial stages of loading when the ion is hot and the laser must be detuned far to the red of the transition. In addition, a strong cooling transition will give a stronger fluorescence signal, thus making it easier to detect, especially in the early part of an experiment when the loading and alignment procedures are still being established.

3.1 Group IIIA Ions

To perform high resolution spectroscopy it is desirable for the clock transition to be field-independent: a transition that is relatively immune to shifts from magnetic fields and from field gradients. Stark shifts due to ambient electric fields for laser cooled single ions are usually quite small and can be ignored. This is because a laser cooled ion will reside very close to the center of the trapping potential. If there is an ambient DC electric field in the vicinity of the trap, it simply serves to shift the position of the potential center to a point where the potential cancels the DC bias. That is, there is effectively no bias. There is, of course, an AC Stark shift which I will come to in a moment. Since the electronic Zeeman shift is proportional to m_J , one way to have a transition free from such shifts is to make sure that m_J is zero in both states of the transition. If J is non-zero, then the ion must be prepared in the $m_J = 0$ state. This requires additional apparatus to prevent optical pumping into $m_J \neq 0$ states. A simpler way to accomplish this is to use states with $J = 0$, thereby guaranteeing that $m_J = 0$. Such states are forbidden (for a $J = 0 - J = 0$ transition,

selection rules require that $\Delta J = \pm 1$) and only proceeds to the extent that IJ mixing is strong (the IJ -coupling approximation is not valid).

One way to get a $J = 0 - J = 0$ transition is to use an ion with two valence electrons in the same quantum state (up to spin state). This gives an integral total spin which in turn will result in an integral total angular momentum J . Regardless of whether the electrons form a singlet or triplet, there will always be a $J = 0$ state present. Such is the case for singly ionized group IIIA atoms. These atoms have a single p electron that, when stripped away, leaves two s electrons (the group IIIB atoms also have 3 electrons but they are grouped as an inner d electron surrounded by 2 s electrons with higher principle quantum number and ionization leaves an s and a d electron). The ground state is a singlet s with $J = 0$. All group IIIA ions have 1P_1 or 3P_1 excited states that could be used for cooling and a 3P_0 metastable state that could be used for spectroscopy. Each of these states are coupled to the same 1S_0 ground state making the electron shelving technique for amplification of the fluorescence from the weak transition possible. An added advantage of the $J = 0 - J = 0$ transition is that neither the ground or excited state has a quadrupole moment which would interact with field gradients. Since field gradients in RF traps can be rather large, this is an important consideration for ion trapping.

3.2 Cooling Transition Considerations: Choosing In^+

Because of the field-independent clock transition, group IIIA ions would appear to be ideal candidates for high resolution spectroscopy, however there are some complications. The electric dipole transition, $^1S_0 - ^1P_1$, is generally in the vacuum ultraviolet (VUV) making it difficult to synthesize and to manipulate. The next best cooling transition is the intercombination line, $^1S_0 - ^3P_1$. This transition requires a spin flip and so is only as strong as the LS coupling which means that it is generally much weaker than a typical electric dipole transition. For heavier ions, LS coupling is stronger and so the $^1S_0 - ^3P_1$ transition is likely to be strongest. This would suggest Tl^+ , the heaviest group IIIA atom, as the best choice. However the $^1S_0 - ^3P_1$ transition in Tl^+ is at 202 nm and so still on the edge of what can be synthesized. The next heaviest group IIIA ion, In^+ , has the $^1S_0 - ^3P_1$ transition at 231 nm and the metastable clock transition, $^1S_0 - ^3P_0$, at 237 nm—both accessible with current laser technology. See Figure 3.1 for the In^+ energy level structure. For these reasons,




Figure 3.1: Indium II level diagram.

we have chosen to start our investigations of trapped, laser cooled, single group IIIA ions with In^+ .

3.2.1 Transition Lifetimes

Both the $In^+ \ ^1S_0 - ^3P_1$ (cooling) and $\ ^1S_0 - ^3P_0$ (clock) transition lifetimes have been measured [Pei94]. Their values are: $\tau_{cool} = 0.44\mu s$ and $\tau_{clock} = 0.14s$. These and other transition values are summarized in appendix A.

3.2.2 Expected Fluorescence Signal From Trapped Indium Ions

In Section 4.10.2 I will derive expressions for the detection efficiency of our apparatus applied to ion clouds and single ions. Equation 4.13 gives the value

$$E_d = 9.2 \times 10^{-7} \quad (3.1)$$

for the detection efficiency in the ion cloud case. Taking this value together with eq. 2.56 for the emission rate, and assuming a buffer gas cooled cloud of ions at room temperature irradiated by a 5 MHz laser, we get

$$R_{max}(T = 300K, \Delta\omega_L = 2\pi \times 5MHz) = 9.48 \times 10^{-3} \frac{counts}{\mu W \times ion \times s}. \quad (3.2)$$

If we assume 50 μW of laser power incident on the ion cloud and 50,000 ions in the trap we get a predicted signal of $R_{max} = 13,000s^{-1}$. This compares well with observed values (to skip ahead for a moment), but the assumed trap capacity is about an order of magnitude less than the capacity predicted in Section 2.3.3. So we still have an order of magnitude floating around—most likely in the trap capacity calculation. One can also see that a room temperature *single* ion would be unobservable, producing an R_{max} of much less than $1s^{-1}$. If we now consider a laser cooled single ion, using the value of 3.9×10^{-4} for the detection efficiency given in eq. 4.14, we have

$$R_{max}(T \simeq 10\mu K, \Delta\omega_L = 2\pi \times 5MHz) = 661 \frac{counts}{\mu W \times ion \times s}. \quad (3.3)$$

With a laser stabilized below the natural width this may not improve very much due to saturation effects (see the next section). While the counts per μW would go up, the saturation power goes down by roughly the same factor. See Appendix A for a summary of these estimates.

3.2.3 Laser Power Requirements

A general discussion of saturation power was given in Section 2.6. Substituting values for the cooling and clock transitions into eq. 2.59 and assuming a laser width of 5 MHz gives

$$P_{sat}^{cool}(T = 300K) = 402\mu W, \quad (3.4)$$

$$P_{sat}^{cool}(T \simeq 10\mu K) = 1.34\mu W, \quad (3.5)$$

and

$$P_{sat}^{clock}(T \simeq 10\mu K) = 1.24\mu W. \quad (3.6)$$

The last value may seem high for a weak transition, but the laser is quite broad. If we now look at saturation powers with a laser stabilized to 30 kHz we get

$$P_{sat}^{cool}(T \simeq 10\mu K, \Delta\omega_L = 30kHz) = 97nW, \quad (3.7)$$

and

$$P_{sat}^{clock}(T \simeq 10\mu K, \Delta\omega_L = 30kHz) = 7.4nW. \quad (3.8)$$

Finally, consider the saturation power for an ion which is initially at 1/10 the well depth of our Paul trap (about 30,000 K, see Section 4.9) and a 5 MHz laser:

$$P_{sat}^{cool}(T = 30,000K, \Delta\omega_L = 5MHz) = 4mW. \quad (3.9)$$

These values, particularly the last one, suggest that high power for the cooling laser may be needed (since the 231 nm light must be generated by SHG, anything in the mW region is considered high power). On the other hand, 1 μW from the clock laser should be sufficient. This is partly due to the fact that the clock transition is so weak, but also because we never intend to look at it unless the ion is already cold and the Doppler width has been eliminated. See Appendix A for a summary of these and other estimates.

3.3 Hyperfine Structure and Zeeman Shifts

While the $J = 0 - J = 0$ transition of In^+ is immune to Zeeman shifts in the fine structure, this is not true of the hyperfine structure. Here the nuclear spin of 9/2

results in a non-zero value of F and so a potential shift proportional to m_F . With the values tabulated in appendix A and using the following definitions:

$$g_F = g_J \frac{F(F+1) + J(J+1) - I(I+1)}{2F(F+1)} - g'_I \frac{F(F+1) - J(J+1) + I(I+1)}{2F(F+1)} \quad (3.10)$$

$$g'_I = g_I \frac{\mu_N}{\mu_B} \quad (3.11)$$

$$\mu_N = \frac{e\hbar}{2M_P} = \frac{\mu_B}{1.836 \times 10^3} = 5.051 \times 10^{-27} JT^{-1} \quad (3.12)$$

$$\mu_B = 9.274 \times 10^{-24} JT^{-1} \quad (3.13)$$

$$g_J = \frac{J(J+1) + L(L+1) - S(S+1)}{2J(J+1)} + g_S \frac{J(J+1) - L(L+1) + S(S+1)}{2J(J+1)} \quad (3.14)$$

$$g_S \simeq 2 \quad (3.15)$$

$$g_I = \frac{\mu_I}{\mu_N} \frac{1}{I} \quad (3.16)$$

we can determine the various Zeeman shifts. The relevant g values are:

$$g_J(^3P_1) = 3/2$$

$$g'_{I=9/2} = 6.668 \times 10^{-4}$$

$$g_{I=9/2} = 1.224$$

$$g_{F=11/2}(^3P_1) = 0.273$$

$$g_{F=9/2}(^3P_1) = 0.0606$$

$$g_{F=7/2}(^3P_1) = -0.333$$

There are three hyperfine lines associated with the cooling transition. In first order perturbation theory, the hyperfine splitting is given by ([Woo80, p. 174]):

$$\Delta E_{HF} = \langle \gamma I J F M_F | A(J) \mathbf{I} \cdot \mathbf{J} | \gamma I J F M_F \rangle \quad (3.17)$$

where γ specifies the configuration and $A(J)$ is the hyperfine structure constant. Expanding $\mathbf{I} \cdot \mathbf{J}$ gives:

$$\Delta E_{HF} = \frac{1}{2} A(J) [F(F+1) - J(J+1) - I(I+1)]. \quad (3.18)$$

The value of $A(J)$ can be calculated with reasonable accuracy for atoms with one valence electron [Ram53, p. 10], but for atoms with two valence electrons (such as Indium II) the situation is more complicated. Kopfermann [Kop58, p. 139] gives an approach for this case which assumes that the contribution to the field at the nucleus due to the excited electron is small compared to that of the unexcited electron. An approximation to the hyperfine structure constant can be obtained by calculating the constant as if the unexcited electron were the only valence electron. Using this approximation,

$$\Delta E_{HF}^{I,J} \simeq \Delta E_{HF}^{I,s_1} = a_{s_1} I s_1 \overline{\cos(\mathbf{I}, \mathbf{s}_1)} \quad (3.19)$$

where a_{s_1} is the hyperfine constant for just one valence electron and the argument to the cosine term is referring to the angle between \mathbf{I} and \mathbf{s}_1 in the vector model. With

$$\overline{\cos(\mathbf{I}, \mathbf{s}_1)} = \cos(\mathbf{I}, \mathbf{J}) \cos(\mathbf{J}, \mathbf{S}) \cos(\mathbf{S}, \mathbf{s}_1) = 1/2 \quad (3.20)$$

we get:

$$A(J) \approx a_{s_1} = \frac{8}{3} \frac{hcR_\infty \alpha^2 Z^3}{n^3} g'_I = 1.95 \times 10^{-22} J. \quad (3.21)$$

This yields

$$\Delta E_{HF}^{I,J} \simeq \frac{9}{8} a_{s_1} = 2.20 \times 10^{-22} J \quad (3.22)$$

and corresponds to a frequency shift of $\Delta\nu = 331 \text{ GHz}$. Unfortunately this is about an order of magnitude too high (based on observations)¹ so the approximation must not be appropriate for In^+ .

The hyperfine structure constant is obtained from observations by using the interval rule [Woo80, p. 175]:

$$\Delta E(F) - \Delta E(F - 1) = A(J)F. \quad (3.23)$$

In our In^+ work we observed a difference in wave numbers² between the $F = 11/2$ and $F = 9/2$ states of 1.30 cm^{-1} which corresponds to a frequency shift of 39 GHz

¹ Measurements by Peik et al. [Pei94], Paschen et al. [Pas38], Larkins et al. [Lar93] and our own work are in very close agreement on the hyperfine splitting.

² The error in the device measuring the frequency was in the last digit, but the measurement was only made once so I don't know what other errors may be present, and I don't attempt to determine error bars.

and an energy difference of $2.59 \times 10^{-23} J$. From this we get a value for the hyperfine structure constant of:

$$A(J) = 4.70 \times 10^{-24} J = .237 \text{cm}^{-1}. \quad (3.24)$$

If the ambient magnetic field is not zero there will be a Zeeman shift for each of the magnetic sublevels. In the weak field regime where $A(J) \gg g_J \mu_B B$ (for Indium the weak field regime corresponds to $B < 0.1 \text{ T}$):

$$J \neq 0: \quad \Delta E_{F,M_F} = \frac{1}{2} A(J) K + g_F \mu_B B M_F \quad (3.25)$$

$$J = 0: \quad \Delta E_{F,M_F} = g'_I \mu_B B M_I \quad (3.26)$$

where $K = F(F+1) - I(I+1) - J(J+1)$. From this we get a Zeeman splitting in the excited state of:

$$\Delta E_Z^{e_s} = g_F \mu_B B \quad (3.27)$$

and in the ground state of:

$$\Delta E_Z^{g_s} = g'_I \mu_B B. \quad (3.28)$$

A 1 Gauss field gives $\Delta E_Z^{e_s} = 2.53 \times 10^{-28} J$ or $\Delta \nu_Z^{e_s} = 382 \text{kHz}$ and $\Delta E_Z^{g_s} = 6.18 \times 10^{-31} J$ or $\Delta \nu_Z^{g_s} = 932 \text{Hz}$. The ground state shift is the same as that for the 3P_0 excited state since $J = 0$ in both cases. So the Zeeman shift for the excited state in the clock transition is roughly 400 times smaller than that of the excited state in the cooling transition. At about 1 kHz per Gauss it can be reduced to the level of a part in 10^{18} by shielding the ambient magnetic field to the μGauss level.

3.4 The Stark Shift

The Stark shift for trapped ions is problematic. In non-degenerate perturbation theory the Stark shift doesn't come in until second order and is usually small. This is particularly true, as has already been pointed out, in the case of trapped ions because the ion positions itself at the point where the fields cancel out. However if there is a small DC bias electric field present then the ion is pushed to a region where the pseudopotential is non-zero (so that it can cancel the DC bias). Since the pseudopotential is being created by an RF field, the ion will now experience a non-zero RF field even if it is in the zero point energy level of the pseudopotential. Put

another way, the time-average of the field amplitudes experienced by the ion cancel to zero, but the time average of the *squares* of these amplitudes do not. The quadratic Stark shift is proportional to E^2 so the ion will now experience an AC Stark shift because the time average value of E^2 is non-zero. Whereas the original DC bias might have been quite small (on the order of a volt), the resultant time average value of E^2 can be quite large because the amplitude of the applied RF is usually quite large (on the order of hundreds of volts or more). Thus, if no DC bias is present, then the AC Stark shift due to the trapping field will be zero. However, even if there is a small DC bias, the Stark shift due to this bias should be small, but it will create an AC Stark shift which can be quite large. Consequently DC biases must be avoided and this is one reason why the heatable Paul-Straubel traps, which allow one to remove DC-bias-creating impurities, are desirable.

3.5 Isotopes

There are only two stable isotopes of Indium: $A = 113$ which has a 4.3% natural abundance and is completely stable; and $A = 115$ which has a natural abundance of 95.7% and a lifetime of about 10^{14} years (see appendix A and [Rag89]). Both isotopes have a nuclear spin of $9/2$ which makes it unnecessary in our initial work³ to determine which isotope has been trapped—both have the same level structure and characteristics. The slightly different masses lead to slightly different trap secular frequencies at about the 1 percent level.

³ Although Zhao et al. [Zha95] have used the slight isotope induced change in the secular frequency to selectively load different isotopes and measure the isotope shift in Ba^+ ions.

Chapter 4

EXPERIMENTAL APPARATUS

In this Chapter I will describe the apparatus used to trap single Indium ions. It turned out that this was very similar to the apparatus used to trap large buffer gas cooled clouds. The latter differed from the former primarily in that it had a mechanism for introducing Helium buffer gas into the vacuum chamber and it used a conventional RF trap with endcaps. The single ion trap used was a Paul-Straubel trap.

4.1 Overview

See Figure 4.1 for a schematic of the apparatus. The light source consists of a dye laser pumped by an ion laser. The dye laser output at 462 nm passes through beam steering optics and then into a beam splitter. The beam splitter picks off a small portion of the beam for locking and stabilization purposes. The first part of this is a Tellurium saturation cell for long term stability. The laser is locked to a Tellurium line and tuned relative to that line via the frequency of the voltage applied to an AOM. Another piece of the picked off beam is sent to the short term stabilization apparatus. Here the beam is frequency modulated and reflected off of a high finesse cavity. The result is an error signal which can be used to stabilize the laser well below the level of stability provided by the commercial cavity supplied with the laser. After passing through the beam splitter, the primary beam then enters the Second Harmonic Generation (SHG) Enhancement Cavity (EC). In the cavity is a non linear crystal which doubles the fundamental into the second harmonic at 231 nm. The fundamental that escapes through the output coupler is blocked while the UV is steered towards the trap vacuum chamber. It is focused into the trap with a UV transmissive corrected lens. The vacuum chamber consists of a quartz tube with quartz windows fritted on. Connected to the tube are a triode ion pump and a Helium leak. The latter can be filled with Helium and pumped out whenever it is needed. Fluorescence from trapped ions is then transmitted perpendicular to the




Figure 4.1: Experimental Apparatus.

laser beam into a spherical mirror. The mirror reflects the fluorescence back through a correction plate and into the Photo-Multiplier Tube (PMT) where the light is collected and transmitted to the computer. The correction plate in front of the PMT is used to correct for astigmatism in the output beam.

4.2 Laser Sources

The laser used for the In^+ cooling transition is a Coherent 699-21 ring dye laser pumped by a Coherent Innova 400-25 Argon ion laser operating in multiline UV mode (for maximum power). The pump laser is fitted with UV optics. When it was new, the pump laser produced 7.7 Watts CW with its power supply set to a maximum of 60 Amps. After several months this was reduced to roughly 6 Watts where it tended to stabilize. Initially a Coherent CR-12 ion laser was used as a pump but it generally did not produce more than about 3.5 Watts. This was marginal at best in the early stages of the experiment when maximum power was needed. We were fortunate to have the use of the Innova 400 which was loaned to us from the University of Washington Physics Department's Laser Facility. In addition to increased power, the Innova 400 has a servo mechanism to keep the power output steady ("Power Track") making it largely a turn-key system.

The pump laser drives the 699 dye laser which uses Stilbene 3 (also known as Stilbene 420). With new dye and 5 W of pump power this laser produces up to 500 mW single frequency at 462 nm. With a pump power of 3.4 W the dye laser occasionally produced as much as 300 mW. This quickly degrades however since the dye has a lifetime of about 100 W-hours. It was not uncommon for us to run the laser all day with 5 W of pump power. During these times it was necessary to change the dye every few days.

4.2.1 Dye Change Procedure

The dye recipe used was 750 mg of Stilbene 3 dye dissolved in 1 l of Ethylene Glycol (EG). The procedure for mixing the dye was as follows. The dye, in powdered form, was put in a small beaker and 100 ml of EG was added. This was placed on a magnetic stirrer with low heat until no undissolved dye remained. Meanwhile the old dye in the 699 dye recirculator is emptied by detaching the hose at the jet assembly

input and placing it into a bottle while running the recirculator. To get as much old dye out as possible, the recirculator pressure is turned up to maximum. The hose is then reattached to the jet assembly and residual dye in the recirculator dye reservoir is poured out. The recirculator is then turned upside down on top of a bucket and the filter housing is removed, taking care not to lose the filter O-ring and to keep track of its orientation with respect to the filter. The filter and O-ring are removed and the housing reinstalled. The recirculator is taken off the bucket and the dye reservoir is filled with Methanol. The recirculator is then turned on for several minutes to allow the Methanol to clean out any remaining dye. The dye hose is again disconnected and placed into a waste can and the recirculator is emptied as before. With the recirculator again upside down the filter housing is removed and a new filter (with O-ring correctly oriented) is installed. We found that it was essential for the O-ring to be oriented and seated correctly. If either of these aren't done correctly, bubbles in the dye jet will result leading to large noise in the laser power. To best seat the O-ring, I would first install it, along with the filter and filter housing, while the recirculator was upside down and then tighten the filter housing further when the recirculator was right side up and I was able to get better leverage.

Finally the remaining 900 ml of EG is added to the dye reservoir and the recirculator is turned on. The jet pressure is set to approximately 40 psi. Allowing the EG to circulate for several minutes will filter out many of the air bubbles. Once the dye is completely dissolved on the stirrer I would usually place the beaker in a pan of cold water and continue to stir (with the heat off) to bring the dye to room temperature. The dye mixture is then slowly added to the dye reservoir. This is given several minutes to mix in, at which point the dye change is complete.

4.2.2 Reacquisition of Lasing after Dye Change

In order to maintain high 699 power output, as the dye ages the pump laser power is usually increased. Consequently, after a dye change the pump power is often higher than necessary. When the pump power is increased, maximal 699 power is achieved by defocusing the pump beam somewhat (moving the focal point further in front of the dye jet). This is to prevent the thermal lensing effect whereby if the pump beam is too strongly focused in the jet, thermal effects will destroy the quality of the beam and result in overall power loss. After a dye change, the pump power is

brought back down and the pump focus must be brought back in. Usually the 699 should lase immediately with little effort after a dye change. If it doesn't, something is probably wrong. One possibility is a dirty Intra-Cavity Assembly (ICA). We had great difficulty for a period of time for this reason. Any change in one of the laser cavity optics, such as changing the pump focus, would completely throw the laser out of alignment. Presumably the thick etalon inside the ICA had "hot spots" where there was no dirt. It took a great deal of effort to get the laser aligned so that power was optimal and the beam was going through a hot spot. Any change in an optical element would move the beam and reduce power dramatically. Finally we had the ICA etalons cleaned by Coherent and the problem completely disappeared.

Another potential impediment to laser oscillation action can be incorrect dye mixture which can lead to the wrong amount of pump beam absorption in the dye jet. If one calls Coherent for technical assistance on a dye laser one of the first things they will ask is whether the dye absorption has been checked. I got used to checking it before I would call them and set up a pump laser high reflector (HR) on a thin post that could be inserted into the 699 cavity to check absorption. The procedure is as follows. First, turn the pump power way down to less than 10 mW. With the dye jet off and the dye splatter shield retracted, the weak pump beam is reflected off of the pump mirror through the dye jet area and onto the inserted HR. The HR then steers the beam to a power meter. With the beam thus safely directed, the pump power can now be turned back up to a normal value. The power is measured with and without the dye circulating to find how much power is being absorbed. A value of 75-80% is considered optimal by Coherent.

4.2.3 Jet Assembly Problems after Dye Change

If the 699 lases with correct power (250 mW at 462 nm with a pump power of 3.4 W was considered normal for our Stilbene 3 system) after the dye change but is noisy, it is likely that there are air bubbles in the dye mixture. These will generally be cleared out by the filter after some time circulating. If they don't, then it is likely that some matter has gotten stuck in the dye jet nozzle. At this point two steps need to be taken. First, an error in the seating of the filter is the likely cause for the contamination of the nozzle so this must be checked and corrected. Second, to remove the contamination, the nozzle must be cleaned. If the nozzle itself is removed to do so

this may result in considerable realignment effort when it is reinstalled. Consequently, to clean the nozzle it is much better to remove the entire jet assembly which mates to the cavity on 3 ball bearings. These ball bearings completely determine the position of the jet assembly within the cavity and make reinstallation trivial. To remove the jet assembly, retract the dye splatter shield, slide the rubber “boot” off the jet assembly, remove the dye hose and loosen the screws just below the connection of the dye hose to the jet assembly. The assembly will just lift out. I typically cleaned it by placing it in a large beaker filled with Methanol and suspending the beaker in an ultrasonic cleaner for about 10 minutes. Care must be taken not to let the jet nozzle touch anything during the cleaning process as this could damage it.

4.3 Laser Stabilization

“When working with RF, if you design an amplifier you get an oscillator and if you design an oscillator you get an amplifier.” — *Fritz Toevs*

In order to do high resolution spectroscopy a frequency stable clock laser is required. This has not been done yet on the In^+ experiment. Though not required in principle, we found it advantageous to stabilize the cooling laser. There are two reasons for this. First, a stable cooling laser has a narrower line width and therefore a higher spectral density. This makes it possible to get the same amount of ion signal with less laser power. Less laser power means lower background and a better signal to noise ratio. The commercial stabilization for the 699 laser supplied by Coherent gives a laser line width of 1 MHz RMS at best¹. By locking the cooling laser to a Fabry-Perot Interferometer (FPI) we were able to achieve a laser line width of approximately 30 kHz, so it is possible to get an order of magnitude decrease in background with no loss in signal. The second reason for stabilizing the cooling laser is that it becomes possible to tune the laser closer to, but still below the ion transition frequency without heating the ion. Being able to tune the laser closer to resonance means an increase in signal with no increase in background. Both of these factors taken together make

¹ In our experience, the laser rarely did better than 5 MHz as measured against a Te_2 saturation line. The 1 MHz value is determined by observing RMS fluctuations in the error signal. It doesn't take into account cavity drift or high frequency components which may lead to a broader line.

it easier to see the weak cooling transition *and* make it possible to observe quantum jumps using only this transition (see Section 2.7).

There are two sides to stabilization. First there is high frequency “jitter”. The most likely cause for jitter in dye lasers is the presence of air bubbles in the dye stream. A dye circulation speed of 1 m/s and an average bubble size of 1 μm will give frequency excursions of roughly 1 MHz. The width of the In^+ cooling transition is 360 kHz. To resolve this natural width the laser must be narrower than 360 kHz. Ultimately, high resolution spectroscopy is performed on the “clock” transition. For In^+ the clock transition width is on the order of 1 Hz ($1/2\pi\tau$). Obviously the limiting factor in the precision of this spectroscopy is determined by the stability of the clock laser. Jitter is greatly reduced by locking the laser to a high finesse cavity using an FM sideband technique. Hereafter I will refer to the stabilization of short term hops, or jitter, in the laser frequency as *short term stabilization*.

The second stability issue is long term drift. Measurements on our 699 showed that it drifts roughly 1 MHz/min. We would usually scan the laser frequency at 1 sec/channel where each channel would cover about 0.4 MHz. Even without the laser stabilized for jitter the cooling transition linewidth was in the 10 MHz range which, at our typical scan rate, would take about 25 s to traverse. In this amount of time the laser drift is appreciable and effects the signal. Perhaps more important is the effect that drift has on optical optimization. To get the best possible signal, one “parks” the laser as close as possible to the ion transition without laser heating and then tunes the optics for maximum signal strength. This is virtually impossible to do with a drift of 1 MHz/min on a laser cooled single In^+ ion because the signal strength is changing with the drift. To eliminate the drift the laser is locked to a molecular line obtained from a Tellurium saturation cell. Just as important as eliminating drift is the ability to roughly locate the laser in the vicinity of the ion transitions of interest. This can be accomplished (expensively) with a wavemeter, however we were able to do quite well simply using the Tellurium spectrum. The pattern of Tellurium lines near the ion transition is unique and relatively easy to find. At the start of each day it usually only took a few minutes of searching the Tellurium spectrum to find the desired laser frequency. Hereafter I will refer to elimination of laser drift as *drift stabilization*.




Figure 4.2: Short Term Stabilization Apparatus

4.4 Short Term Stabilization

The FM sideband technique used to reduce the short term frequency excursions or “jitter” is described in many places (see for example, [Dre83]). Here I will restrict the discussion to our implementation of it; the optics and the electronics used for locking.

4.4.1 Apparatus

See Figure 4.2 for a diagram of the apparatus used for locking. A small amount of the primary beam is split off and passed through an electrooptic modulator, or EOM, (ConOptics model 370 [EOMCon]) which has a drive frequency of 23.5 MHz. The optical axis of the EOM is positioned such that the beam is frequency modulated (optic axis parallel to incident polarization) with sidebands at plus and minus 23.5 MHz from the laser frequency. A step-up coil is placed at the input of the EOM to give an RF amplitude of approximately 100 V. The frequency chosen is based on the availability of crystal oscillators and the desirability of having the carrier be substantially larger than the frequency of the error signal (typically as much as 5 MHz).

The laser generates linearly polarized light in the vertical direction so this is incident on the stabilization optics. The light passes through a polarization beam splitter oriented to transmit this polarization. Next the beam passes through a quarter wave plate creating circular polarized light. This is now incident on our reference cavity. For the cooling laser stabilization we found it sufficient to use as our reference cavity a Tropel model 240 spectrum analyzer (Fabry-Perot interferometer). This cavity has a mirror separation of about 5 cm giving a free spectral range of about 3 GHz. The finesse of the cavity is 3000 so that the cavity width is about 1 MHz. Tropel builds this cavity in the confocal configuration. If m and l represent the transverse mode numbers, n is the longitudinal mode number, and L is the cavity length then for a confocal cavity, modes with the same value of $2n + m + l$ will have degenerate frequencies [Sve89, p. 174]. The best signal is achieved when the beam is mode matched into the cavity and only the TEM_{00} mode is resonant. To resolve all of the modes separately, making it possible to determine when the cavity is correctly mode matched, the degeneracy just mentioned must be removed. This is accomplished by moving the mirrors just slightly away from their confocal arrangement.

Some amount of the circularly polarized beam is reflected off the cavity input coupler (call this the input beam) and some amount leaks out of the input coupler after being transmitted into the cavity (call this the output beam). These two beams pass through the quarter wave plate which now produces linear polarization orthogonal to the original polarization and they are reflected by the polarization beam splitter. The reflected beams are detected with a high frequency photodiode. The sidebands on the input beam have opposite signs for frequency modulation. Only the part of this beam resonant with the cavity will leak back out as part of the output beam. In addition, the output beam has a 180° phase shift (negative sign). When the laser is right on the cavity resonance, the carrier will be cancelled by the addition of the input and output beams leaving only the sidebands in the reflected input beam. This gives a DC signal at the detector. If the laser is slightly off resonance, the amplitude of the carrier in the output beam will be reduced and when it is added to the input beam carrier it will not be completely cancelled. The signal at the detector will be a combination of the sideband and carrier strengths plus two beat note components with opposite signs and unequal amplitudes. The beat notes are produced by the sidebands beating with the carrier which is only possible off resonance. As the laser shifts frequency up

or down from resonance one sideband will be weighted more heavily than the other and will contribute more positively or more negatively to the overall signal. That is, while the resonance lineshape doesn't change sign as the laser sweeps through the resonance frequency, the detected signal does. Hence, as the laser shifts up or down in frequency, the detected signal will follow a dispersion lineshape ideal for use as an error signal.

By expanding the frequency modulated laser light in Bessel functions one can show that a signal proportional to the error signal will exist at the modulation frequency, 23.5 MHz [Hal81, Hil87]. After amplification, the detected signal is mixed with the 23.5 MHz drive in a Mini Circuits SRA-1W double balanced mixer. The signal is now being detected far away from the DC laser noise and, in general, the 1/f noise is much smaller at the modulation frequency than at DC. Since the signal is proportional to the amount of power in the sidebands it is necessary to maximize this power. It can be shown that sideband power is approximately proportional to the product of J_0 and J_1 . This product is maximized at a modulation index, $\beta \approx 0.8$. For small β ,

$$J_0(\beta) \approx 1 - \frac{\beta^2}{4} \quad (4.1)$$

and

$$J_1^2(\beta) \approx 1 - \frac{\beta^2}{2}. \quad (4.2)$$

Since the detector sees intensity, we can determine that $\beta = .8$ when the carrier intensity is reduced to 68% of its unmodulated value. Therefore, it is possible to maximize the detectable signal simply by changing the amplitude of the RF voltage applied to the EOM until the carrier is reduced by 32%. The amplitude is changed by adjusting the position of the input line on the step-up coil attached to the EOM.

The signal is electronically integrated, amplified, and then fed back to the laser. There are several transducers in the laser cavity driven by the stabilization electronics. First, there is the galvo driven Brewster plate ("woofer") and the PZT driven High Reflector ("tweeter"). These are supplied by Coherent as part of the 699 and are driven by the 699-21 servo loop. In addition we have inserted an intracavity EOM (Gsänger model PM 25 [EOMGsa]) with very low loss, high transmission windows tilted at Brewster's angle. The two windows are oppositely tilted so that the beam suffers relatively little transverse offset. This works well enough that the EOM can actually be inserted by hand into the laser cavity without loss of laser oscillation.

These three transducers taken together provide a hierarchy in bandwidth. The bandwidths are about 100 Hz for the Brewster plate, tens of kHz for the tweeter and up to a GHz for the EOM (though the electronics bandwidth limits the latter to about 1 MHz). They are operated in parallel by the electronics so at any given frequency the transducer with the most gain dominates. As described in the next section on the electronics, a clever technique takes advantage of the commercially supplied servo loop for low frequency excursions and as a backup for the added high frequency loop. If the high frequency loop loses lock, the circuit will automatically revert to the low frequency loop until the high frequency loop reacquires lock.

4.4.2 Electronics

The electronics used for short term stabilization consists of several parts: error signal detection, EOM drive, slow path transfer, and redundant cavity locking. The overall design is based on the locking circuit described by Sandberg [San93]. The description here is not exhaustive, rather I will focus on areas either that Sandberg did not cover or that we did differently. The reader is referred to Sandberg's thesis for more details. See Appendix C.3 for schematics of our implementation.

The detection electronics consists of an RF photodiode, a low noise preamplifier, and a power amplifier. This produces a signal which is mixed with 23.5 MHz provided by a TTL oscillator. The result is a signal proportional to the error signal present on the sideband. This error signal is fed into the EOM drive electronics which essentially consists of two integrators. One of the integrators has high bandwidth and the other produces high voltage (HV). Since their outputs are not cascaded, but rather applied to either side of the EOM crystal (the EOM case has two inputs) we get the advantages of both without a two-pole rolloff. In addition to being applied to the EOM, the output of the high bandwidth integrator is fed to the slow path transfer electronics, which is used to match the gain of the high and low frequency transducers at the crossover point. This in turn is fed to the redundant locking electronics, which is used to drive the 699 servo loop in a clever way to be described.

The whole circuit, except the detector, is fit into a single Nuclear Instrument Module (NIM)² and runs off of $\pm 12V$ and $\pm 24V$. We decided to make NIM our

²Sadly, this nicely compact configuration didn't work very well. It was necessary to pull the oscillator out of this box and run it externally to avoid excessive pickup.

standard for several reasons. First, once a NIM crate is located one does not need to deal with the perennial problem of setting up separate dedicated power supplies for every new piece of electronic equipment. This not only reduces clutter but simplifies adding new electronics. The NIM crate will hold up to 12 separate modules all operating off of the same set of internal supplies. Second, NIM power supplies are generally quite stable. Third, the NIM crate is generally more space efficient than regular rack mounted boxes since electronics boxes rarely are large enough to really require the whole width of a rack and can usually be fit quite comfortably into a single NIM slot. When this is not true, double slot boxes or multiple boxes can be used.

Detection of the error signal coming from the FPI is accomplished using an EG&G model FND-100Q fast photodiode. This diode has an active area of about 5 mm. making alignment easy, and a bandwidth of 350 MHz into 50 ohms. The latter is needed because our signal carrier is 23.5 MHz. Photodiodes can be operated in either of two modes: photoconductive or photovoltaic. In photoconductive mode the diode is reverse biased to give low junction capacitance:

$$C_j = \frac{1.92 \times 10^4 A_d}{\sqrt{\rho(V_0 + V_{rb})}} \quad (4.3)$$

where A_d is the active detection area of the diode, ρ is the resistivity (ohm-cm) of the depletion layer, V_0 is the contact voltage, and V_{rb} is the reverse bias voltage. A higher reverse bias voltage results in a lower junction capacitance and higher frequency response. In photovoltaic mode the diode is operated with no bias into a resistive load. This configuration is good for detecting weak signals. We use photoconductive mode because our signal is strong and fast.

When designing a photodetector circuit it is essential to be shot noise limited. That is, shot noise is unavoidable whereas other sources of noise, such as Johnson noise, are dependent on the circuit. Ideally, all other sources of noise will be less than the shot noise. For our detector the typical amount of incident light, P , is about 5 mW at 462 nm. The responsivity, \mathfrak{R} , of the FND-100Q is about 0.25 A/W at 460 nm. Thus the signal current is $i_s = \mathfrak{R}P = 1.25mA$ and the associated signal voltage across a $1k\Omega$ resistor is $V_s = i_s R = 1.25V$. The shot noise current is given by the equation [Hor93, p. 432]:

$$I_{SH}(rms) = \sqrt{2qi_s B} \quad (4.4)$$

where q is the electron charge and B is the measurement bandwidth—about 5 MHz in our case. Putting our value for i_s in we get $I_{SH} = 45$ nA. Therefore the shot noise voltage across our $1k\Omega$ resistor is $V_{SH} = I_{SH}R = 45\mu V$. Now the Johnson noise voltage is given by [Hor93, p. 433]:

$$V_J = \sqrt{4kTRB} \quad (4.5)$$

where k is Boltzman's constant, T is the temperature, R is the resistance and B is again the measurement bandwidth. Substituting the relevant values gives, $V_J = 9.1\mu V$. In other words $V_J \ll V_{SH}$ and we are shot noise limited as desired.

The preamplifier used is a CLC (Comlinear Corporation) 425 which has a very high gain bandwidth product (GBP) of 1.7 GHz, low input noise voltage of 1 nV/ \sqrt{Hz} , and a reasonable slew rate of 350 V/ μs . The 425 is configured as a transimpedance amplifier with the gain being determined by the feedback resistor. The signal appearing at its output is simply $V_o = -i_s R_f$. A larger R_f gives more gain but at the price of increased Johnson noise (see eq. 4.5) and reduced bandwidth. The feedback values of $C_f = 2pF$ and $R_f = 1k\Omega$ give excellent Johnson noise and an adequate rolloff frequency of about 100 MHz. Another aspect of choosing the correct feedback values is a Comlinear recommendation that $C_f = C_i$ where C_i is the capacitance of the diode. EG&G indicates that the diode has 8.5 pF capacitance at 90 V bias. Since the capacitance is inversely proportional to the square root of the bias voltage (eq. 4.3), the capacitance for our bias of 15 V is roughly 17 pF. In order to make $C_f = 17pF$ we would have to lower R_f below that required for gain purposes. We chose to use a lower value of C_f which seemed to work well. A signal is taken directly from the output of the CLC 425 as the DC output. It is also AC coupled into a Mini Circuits Mar-3 linear amplifier for power amplification. The output of the linear amplifier is capacitively coupled to the AC output of the detector box.

The output of the detector is connected to a Mini-Circuits SRA-1W double balanced mixer as the RF. The LO is taken from the 23.5 MHz oscillator that is used to frequency modulate the laser. The result at the IF is a DC signal proportional to the error signal that exists on the first sideband. The oscillator consists of a crystal TTL oscillator followed by a low pass Butterworth " π " filter. The filter has 3 dB attenuation at 40 MHz and roughly 50 dB attenuation at 50 MHz (second harmonic). Thus only the fundamental at 23.5 MHz gets through and it produces a nice sine

wave. The 23.5 MHz signal is passed through a phase shifter (which can simply be a length of cable) before it is sent to the mixer. The output of the mixer (DC error signal) is fed into the EOM drive.

The next part of the electronics is the EOM drive. This is really the heart of the circuit as it consists of the integrators which apply the error signal back to the EOM in the laser cavity. The approach is to take advantage of the two inputs to the EOM, one for each side of the crystal. One side is connected to a wideband low gain integrator and the other side is connected to a lower frequency high voltage integrator. In this configuration, the advantages of both can be used without the usual two pole rolloff associated with two cascaded integrators. For the large bandwidth integrator we use another CLC 425. Sandberg [San93] uses an Apex WA01 which has a much higher slew rate (about a factor of 10) than the 425, but the latter is less expensive and seems to perform reasonably well. However, this is one place where the circuit might be improved in the future—especially if it were going to be used in a situation that required much higher frequency correction. For the high gain integrator we use the Apex PA85 which has a 450 V swing along with 1 MHz GBP. This op amp has a bipolar power supply for better frequency response. The feedback capacitor and resistor on the PA85 are set to give unity gain at about 1 MHz and a rolloff at several kHz. At this point the CLC 425 takes over since it has appreciable gain beyond 100 MHz.

After being integrated by the 425, the signal is split into several parts. The first goes directly to the laser EOM as the high frequency EOM drive, the second goes to the “slow path transfer” circuit to be described shortly, and the third goes to the input of the PA 85. In this last connection a potentiometer is placed in series to control the HV gain of this integrator from the front panel. Also a FET switch is inserted that turns the HV op amp on or off depending on whether the laser is locked to the reference cavity or not. The FET is controlled by the “Lock Control” input which is 0 V if the laser is not locked and -15 V if the laser is locked (the generation of the Lock Control value is described in the redundant cavity locking part of the circuit). If Lock Control = 0 V (out of lock) then the NPN transistor hooked up to the gate of the FET will turn on and current will flow through the $10k\Omega$ resistor also connected to the gate. This current will bring the gate low, turning the FET off. The overall effect is to turn off the HV EOM drive when lock is lost. When lock is

reacquired by the redundant locking circuit the HV EOM drive is also turned back on.

The redundant locking electronics is used to make the standard 699 servo loop act as a backup for the high finesse cavity (referred to as the external cavity) based loop. First, if the laser is locked to the external cavity then the error signal fed into the 699 servo loop comes from this cavity. This means that this error signal is now driving *all* of the laser transducers: the intracavity EOM, the tweeter and the Brewster plate. At the same time, the 699-21 is generating another error signal relative to its own cavity. Instead of feeding this back through the 699 servo loop (which is already being driven by the external cavity error signal), this signal is integrated and fed into the 699 external control. This keeps the 699 cavity in resonance with the external cavity. So the laser is not only locked to the external cavity, but both cavities are kept in resonance with each other. The redundant locking electronics determines whether the laser is locked to the external cavity by monitoring leakage light from the back of the cavity. This is referred to as "cavity photodiode in" in the circuit. Now, if the laser is perturbed and loses lock to the external cavity, the HV EOM drive is turned off and the error signal driving the 699 servo loop is switched from the external cavity to the 699 cavity. The 699 is essentially switched to normal operation. This is possible because the reference cavity has been kept close all along by the redundant locking apparatus. As the laser is brought back into lock by the 699 servo loop, it is likely to cross resonance with the external cavity. At this point the cavity photodiode will detect an increased signal and will cause the whole loop to switch to the external cavity for its error signal. Thus, we are able to take advantage of the larger linewidth of the 699 cavity. The larger width means that its error signal has appreciable amplitude further away from resonance. We use this error signal after a laser jump to bring the laser back into the range of the external cavity error signal, at which point it takes over.

As has already been mentioned, the lock control signal is $0V \rightarrow$ system out of lock, and $-15V \rightarrow$ system in lock. Additionally there is a master lock switch on the front panel. In the off (down) position the switch simply holds the lock control signal at ground thereby turning off the HV servo. This will override the lock detection circuit. That is, even if the lock detect photodiode has a strong signal and the lock detect circuit attempts to put $-15V$ on the lock control output, the master lock control will

keep it at ground (which, in turn, keeps the HV servo off). In the on (up) position, the connection to ground is open and as soon as the lock detect circuit sees the lock condition it will be free to turn on the HV servo. Note that turning the front panel switch to *ON* does not imply that the circuit is now or will ever be in lock. It simply indicates that locking is enabled. An LED light directly below the master lock switch will go on as soon as lock is actually acquired.

The slow path transfer electronics is sandwiched between the EOM drive and redundant locking. Its purpose is to supply some tunable gain right at the cross over frequency between the EOM drive and the slower transducers. This happens somewhere near 30 kHz as indicated in [San93]. The additional gain at crossover is supplied to make the transition from the domain of one transducer to another smoother.

4.4.3 Servo Loop Stability

This topic is covered quite well in Sandberg's thesis. A good knowledge of feedback control systems (see ref. [Van86]) and phase locked loops (see ref. [Hor93, p. 641] and ref. [Gar79]) is very useful. I have also found ref. [Sto76] to be quite useful for overall op amp stability questions. Since the 699 servo loop is an integral part of the whole circuit and its transfer function is not known, it is difficult to do a detailed analysis of the entire circuit's frequency response and stability. Nevertheless, it is possible to make some useful qualitative statements. There are essentially 6 elements to consider. The Brewster plate (unity gain at about 100 Hz); the tweeter (unity gain at 10 kHz); the external cavity (unity gain at about 10 MHz); the EOM (unity gain at about 1 GHz); the wideband integrator (unity gain at several hundred MHz); and the HV integrator (unity gain at about 1 MHz). If all of these elements were in series then there would very quickly be an unstable phase shift (6 pole rolloff). Instead, several components are in parallel and when their gain drops below 1 they are effectively out of the circuit with no effect on an overall phase shift. This is true of the Brewster plate and the tweeter. Above 10 kHz only the EOM, the cavity and the two integrators have any effect. As described by Sandberg, the cavity pole at 1 MHz is cancelled by an effective zero due to the flat frequency response of the wideband integrator at this frequency. The EOM rolloff is at a high enough frequency to be of no concern here, leaving just the HV integrator. Hence the whole circuit, which has




Figure 4.3: Apparatus used to lock to a Tellurium molecular line.

very high gain (and multiple pole rolloff) at low frequency, is reduced to one which has just a single pole rolloff at unity gain and is therefore stable.

4.5 Drift Stabilization and Location of Indium Transitions: Tellurium Saturation Cell

See Figure 4.3 for a layout of the apparatus used to lock the laser to a Tellurium molecular line. This apparatus uses saturation spectroscopy [Cor77, p. 414] to achieve a narrow linewidth and the FM sideband technique [Hal81] to derive an error signal. A small amount of the primary laser beam is split off using a microscope slide as a beam splitter. This beam is again split into two paths: one goes through an EOM and then into the Te_2 saturation cell, while the other goes through an Acousto-optic Modulator, or AOM, (ISOMET model 1205C-1) and then into the saturation cell from the other side. Because the second beam is transmitted through the second beam splitter it is the stronger of the two and is referred to as the saturation beam. The other beam is called the probe beam. Because the two beams are oppositely directed they interact with different "Doppler subgroups". That is, because of the

Doppler shift and an atom's velocity, it will interact with the saturation beam at some frequency off resonance. However, the Doppler shift will be opposite in the probe beam and the probe beam will not be resonant. If the probe and saturation beams are not shifted relative to each other then only molecular velocities very close to zero will interact with both beams simultaneously. This narrow band of bi-resonant velocities gives rise to the saturation profile. As the laser frequency is scanned, the saturation beam sweeps out a Doppler broadened lineshape and is given enough power to saturate the transition. When the laser frequency is below the transition the saturation beam is resonant with molecules traveling with some velocity *towards* the beam. Meanwhile the probe beam is also sweeping out a Doppler broadened line, but it is also interacting with molecules traveling towards it which are necessarily different from those interacting with the saturation beam. As the laser frequency reaches the center of the Te_2 resonance, both beams will be resonant with the same group of atoms and because the saturation beam is saturating the transition, the probe beam will generate no fluorescence. Hence, if the fluorescence from the probe beam is monitored a Doppler broadened profile with a narrow dip in it will be observed. The narrow dip is the basis for saturation spectroscopy.

4.5.1 Apparatus

In our apparatus a Te_2 line is chosen which is close to the desired cooling transition and the laser is locked to this transition to prevent it from drifting. In fact, since the pattern of Te_2 lines near this transition is unique, it can be used as a sort of "road map" for finding the correct laser frequency when it is first turned on. We were able to operate in this way without ever needing the use of wave meter except during initial setup.

Of course there is no Te_2 line which exactly overlaps the In^+ cooling transition, so the saturation beam must be shifted relative to the probe beam. In our case we were lucky to have a Te_2 line within 100 MHz of each of the In^+ HF lines that could be used for laser cooling. An AOM is inserted into the saturation beam and an RF signal is applied to it. The In^+ $9/2 - 7/2$ transition lies 68 MHz above the nearest Te_2 line. Therefore, if 68 MHz is applied to the AOM then the second down-tuned (towards the transducer) sideband will be resonant with the Te_2 line at the same time that the unshifted laser is resonant with the In^+ transition. Note that while the second

sideband of the saturation beam is shifted a full 136 MHz from the fundamental, the Doppler subgroup which is simultaneously resonant with both beams at a frequency half way in between. So the second sideband gives the required effective shift of 68 MHz. The 68 MHz RF signal is also mixed with a 10 kHz chopping signal to facilitate phase sensitive detection.

Both the saturation and probe beams are loosely focused by 250 mm focal length lenses. This focal length was chosen to give a beam confocal parameter roughly equal to the saturation cell size. After going through the lens the saturation beam is reflected by a 50% beam splitter which sends it through the saturation cell. On the other side of the cell the probe beam is modulated at 4 MHz by a ConOptics EOM. The modulation of the probe beam gives an error signal as described in [Hal81] that is ideal for locking purposes. After being modulated the probe beam is directed by a high reflector into the saturation cell. The probe and saturation beams are made collinear by overlapping spots on either side of the cell.

The cell consists of a 1 inch diameter Pyrex tube approximately 6 inches long with diagonal windows roughly near Brewster's angle. Several grams of powdered Tellurium are introduced into the cell. A Tungsten coil is wrapped around the cell. Each end of the coil is attached to the two output leads of a Variac which is used to run a current through the coil and heat the cell. The cell and coil are fit snugly between shaped fire brick which in turn is contained in an Aluminum box with holes at either end for the cell windows. The box holes are covered by microscope slides to transmit the light but retard heat loss. A thermocouple is inserted into the box so that it just touches the cell. For optimal signal, the temperature of the cell must be set correctly. Too low of a temperature and there isn't enough vapor pressure of Tellurium. Too high and the signal begins to suffer appreciable collisional broadening. The signal strength was measured as a function of thermocouple voltage (hence temperature) and an ideal setting of 16.1 mV was found. For our setup this corresponded to about 32.5 V on the Variac and a temperature of roughly 400° C (using the 25°/mV rule). Other experiments in our lab had success using a commercially supplied Tellurium cell, but these cells use a single isotope of Tellurium. We tried this, borrowing an existing cell, and found that none of the Tellurium lines were within several hundred MHz of our Indium lines. On a suggestion from Shawn Hartley, an undergraduate working in our lab at the time, we decided to build our own cell with naturally

occurring Tellurium isotope concentrations. This turned out to be very successful. With the natural isotope concentration we found Tellurium lines within 100 MHz of all three hyperfine In^+ cooling transition lines. This made shifting of the saturation beam with an inexpensive AOM quite accessible. In addition, with the help of our expert glass blower, Bob Morely, our homemade cell was substantially less expensive than the commercial variety.

4.5.2 Electronics

On exiting the cell, the probe beam is transmitted through the 50% beam splitter and detected by a photodiode. The signal from the photodiode is amplified and mixed with the 4 MHz oscillator to get a signal on the sideband away from carrier noise. It is then mixed again with the 10 kHz oscillator for phase sensitive detection. The latter is done as usual to detect the error signal away from the laser frequency thereby improving the SNR. See Appendix C.4 for a schematic of this electronics.

During most of the experimental work, the error signal was simply displayed on an oscilloscope. The handy graduate student then completes the feedback loop by manually operating the laser external frequency control and keeping the laser near the desired frequency. This was quite sufficient for the ion cloud work that we did where the linewidths were on the order of GHz (recall the laser drift is only 1 MHz per minute). However, for single ions where the cooling transition natural width of 360 kHz is resolved with a stabilized laser, we needed to eliminate the laser drift. To this end, the saturation cell error signal is fed into a simple single pole integrator. From here it is fed into the external frequency control of the laser. This loop has perhaps only 1 kHz of GBP, but that is more than enough to correct the drift which is quite slow. It is important to note that the absolute position of the Te_2 lines is highly dependent on the temperature of the cell [Mci87], so that temperature stabilization is essential to high precision work. The temperature of the Te_2 cell is not currently stabilized.

4.5.3 Operation of Laser Stabilization Apparatus

First consider the short term locking electronics. As already mentioned, the electronics box fits into a NIM crate and requires $\pm 12V$ and $\pm 24V$ power from the crate. The inputs are:

- **RF ERR IN:** the error signal from the detector/amplifier box.
- **Δ REF IN:** the box generates its own clock signal at 32.5 MHz. This was later replaced with an external 23.5 MHz oscillator. This signal is applied to the modulation EOM and is mixed with RF ERR to produce a DC error signal. Because the modulated RF ERR signal may be phase shifted from REF OUT, Δ REF IN is provided so that the phase of the reference can be shifted before it is mixed with RF ERR IN. In addition REF OUT should be attenuated to 200 mV before being input to Δ REF IN. In the current configuration, REF OUT is not used.
- **699 ERR IN:** this is taken from the 699-21 control box board 1A9 between TP1 and R17. The circuit is broken at this point and routed to a back panel BNC connector.
- **LCK DET IN:** the signal from the cavity leakage photodiode which determines whether the laser is locked to the external cavity.
- **\pm HV IN:** the two high voltage supplies for the HV op amp. Ideally this should be between ± 100 V and ± 200 V, but can be operated in the range: ± 15 V to ± 225 V.

The outputs are:

- **REF OUT:** the internal 32.5 MHz clock. This has been replaced by an external 23.5 MHz oscillator. The output of the oscillator should be connected to the modulation EOM and, after phase shifting and attenuation, to Δ REF IN.
- **699 ERR OUT:** the error signal that is applied to the 699 servo loop in place of its own error signal. This is connected to the ERR IN connector on the 699-21 backpanel that is connected to the other side of the break in the circuit at TP1 (see the description for 699 ERR IN).
- **699 EXT OUT:** the error signal applied to the 699 external frequency control to keep the 699 cavity in resonance with the external cavity.

- **SLOW EOM OUT:** connected to one side of the laser EOM. This is the error signal generated by the HV op amp.
- **FAST EOM OUT:** connected to the other side of the laser EOM. This is the error signal generated by the wide band EOM.
- **LOW VOLTAGE POWER OUT:** on the back of the NIM module is a lab standard ± 15 V and ± 5 V output to be connected to the power input of the detector.

Other controls:

- **SLOW PATH GAIN:** gain adjust for the slow path transfer circuit. Adjust to reduce oscillations at the cross over frequency between the slow and fast transducers.
- **HV GAIN:** gain adjust for the HV op amp.
- **LCK DET OFFSET:** a trim pot on the front panel to set the point at which the lock detection circuit decides that the laser is locked to the external cavity.
- **LOCK SWITCH:** the master lock input. When switched off (down) locking is disabled and the 699 will use it's own cavity for locking. When switched on (up) locking is enabled and the laser will be locked to the external cavity as soon as the circuit is able do so.
- **LCK DET LED:** is lit when the laser is locked to the external cavity.
- **699 LCK DET LED:** is lit when the laser is locked to the 699 cavity (not yet implemented).

Once the laser is locked to the external cavity, its frequency can no longer be tuned by the external frequency control (either the front panel knob or the back panel BNC). It can only be tuned by changing the length of the external cavity. This is accomplished by changing the frequency being applied to the saturation cell AOM. A change in this frequency causes a shift in the sideband thereby changing the lock point

of the saturation cell. If the error signal from the saturation cell is now connected to the external cavity PZT (instead of to the laser external frequency input), then the external cavity will shift its lock point as well. In fact, this is how the whole apparatus should be set up even if changing the laser frequency isn't needed.

The connections to the saturation cell (drift stabilization) electronics are as follows (going from left to right on the front panel).

- 4 MHZ IN: externally derived 4 MHz sine wave. This same signal should also be connected to the 20 dB amplifier and then to the saturation cell EOM. For small phase adjustments use the frequency control on the 4 MHz signal generator.
- 4 MHZ FREQ (pot): frequency of internally generated 4 MHz signal—not currently used.
- PHI (pot): first pot from the left marked PHI. For adjustment of the internally generated 4 MHz phase—not currently used.
- SIGNAL: input from saturation cell photodiode.
- PHI (pot): second pot from the left marked PHI. For adjustment of the internally generated 10 kHz chopping phase (used for phase sensitive detection).
- 180° PHI (switch): unmarked and below the PSD continuous phase adjust pot. Used to shift phase by 180°.
- 4 MHZ (out): output of internal 4 MHz signal. Connected to EOM—not currently used.
- 10 KHZ (out): output of internal 10 kHz chopping signal. This should be connected to the IF port of a double balanced mixer (SRA-1W). The LO port of the mixer is connected to a signal generator capable of going up to 120 MHz. The RF port of the mixer is amplified and then connected to the saturation cell AOM. The result is a roughly 100 MHz chopped signal applied to the AOM and used to tune the lock point of the saturation cell.

- τ (unmarked knob): sets the integration time constant to gain a better SNR at the expense of slower response.
- LEVEL and GAIN (two unmarked pots): the LEVEL adjusts the zero offset of the error signal, which is particularly useful if this signal is to be fed into locking electronics. GAIN is as the name implies.
- OUT: error signal output. Connect either to the saturation cell locking electronics or to the short term stabilization external cavity.

The saturation cell is relatively simple to operate. After the spots on either side of the cell are aligned a strong line is usually visible right away. Connect the output to a scope and scan the laser frequency over a 10 GHz range (almost regardless of where the laser frequency is, a 10 GHz range will include a strong Te_2 line). Once a strong line is located optimize the optics on either side of the cell for maximum signal (it may be necessary to locate a weaker line if the original strong line begins to saturate the electronics).

4.6 Second Harmonic Generation

For the In^+ cooling transition we need a laser beam at 231 nm. Since there are no tunable lasers which produce this wavelength directly we must use frequency synthesis techniques. The dye laser produces radiation at 462 nm which is then doubled to 231 nm using Second Harmonic Generation in a $\beta-BaB_2O_4$ (BBO) crystal. Particularly in the early stages of the experiment, when none of the optical alignment was optimal, it was useful to have a large amount of power in the UV. Since the conversion efficiency in BBO at our wavelength is roughly 10^{-5} , we were getting only a few μW at best. To improve on this we decided to place the BBO crystal inside an enhancement cavity. In the next several subsections I will describe the crystal we used, the enhancement cavity, and the electronics used to lock the enhancement cavity to the laser. The latter makes it possible to tune the laser while still maintaining enhanced UV power.

4.6.1 BBO Crystal

For our first crystal we are indebted to John Sandberg who loaned us one that was used for his Ph.D. thesis project at MIT. This crystal was AR coated for 486 nm and

243 nm. These wavelengths were deemed close enough to ours that the crystal would probably generate reasonable power for us. It was cut for Type I phase matching (both input photons are ordinary rays) and a phase matching angle of 60° . The crystal was positioned on a rotatable mount to facilitate angle tuning. It was held in place on this mount by a small piece of copper bent to apply enough pressure to prevent it from moving but not so much to damage it. We decided that this was a better arrangement than some sort of set screw where damaging pressure would be very easy to apply. With 300 mW of power at 462 nm incident on the crystal we were able to get between 5 and 10 μW of UV out. Later, in the enhancement cavity, we were able to get 100 μW of UV with only 50 mW at 462 nm.

This borrowed crystal did quite well, but we eventually ordered our own. The crystal is placed at a waist in the enhancement cavity with $\omega_0 = 20\mu\text{m}$. This Gaussian beam has a confocal parameter of:

$$b = \frac{2\pi\omega_0^2}{\lambda} = 5.44 \times 10^{-3},$$

where $\lambda = 462\text{nm}$. If the crystal is non-critically phase matched (phase matching angle between the second harmonic and the optic axis is 90°) the optimal crystal length is $5.68 \times b$ [Yar89, p. 401]. However, in our case the crystal was critically phase matched. In addition, this formula doesn't take into account the small degree of angular acceptance in BBO. Angular acceptance is a measure of how much the beam can deviate from its phase matching angle before there is appreciable loss in conversion efficiency. This becomes important for long crystals (more accurately, for situations where the ratio of crystal length to confocal parameter is large) where the Gaussian beam has expanded to the point where its "sides" are propagating at a different angle than its axis. A more sophisticated approach was therefore needed to determine the optimal crystal length. The definitive reference on this subject is a detailed paper by Boyd and Kleinman [Boy68]. This paper covers all of the important general points in great detail, but does not discuss BBO specifically. An approach which covers BBO (and is much easier to read!) is provided by A. Steinbach, M. Rauner, and J. Bergquist [Ste95, Rau94] who systematically measured conversion efficiencies in BBO while varying beam parameters.

In [Rau94] the crystal length is held fixed and Rauner attempts to find the optimal waist. We have the opposite problem. We had already constructed our enhancement

cavity along with its special optics and so our waist size was more or less constrained to $20\mu\text{m}$. In addition Rauner is doing all of his work at a fundamental wavelength of 515 nm and a doubled wavelength of 257 nm. Nevertheless, the indices of refraction differ only slightly between his wavelength and ours. He was able to show that at 90° phase matching where there is no walk-off (see [San93, p. 107] for a nice discussion of walk-off), the optimal length occurs at $l = 2.84 \times b$ where b is the confocal parameter. At the other extreme, where walk-off is large, the optimal length occurs at $l = 1.39 \times b$. Assuming for the moment that we have relatively large walk-off and using our confocal parameter, we get a length of $l \approx 8\text{mm}$. To determine relative walk-off size, Rauner defines the walk-off parameter,

$$B = \rho\sqrt{lk_1}/2 \quad (4.6)$$

where $k_1 = 2\pi/\lambda_1$ is the wave number of the fundamental. A value of $B = 0$ corresponds to no walk-off and a value of $B \geq 10$ is considered large walk-off. A length of 8 mm in our case gives a walk-off parameter of $B = 15$, thereby justifying the assumption of large walk-off. Hence a length of 8 mm was chosen. An aperture, or effective face size of $3\text{ mm} \times 3\text{ mm}$ was chosen as adequate since the beam remains tightly focused over the length of the crystal. I say “effective face size” because the input and output surfaces were Brewster cut so the aperture refers to the cross sectional area perpendicular to the longitudinal faces of the crystal.

When purchasing a doubling crystal one must decide whether to have surfaces AR coated or Brewster cut. Losses due to reflection at the faces for a given wavelength are probably less for a Brewster cut face than for an AR coated face. However, since the surface can only be at Brewster’s angle for one of two orthogonal directions at a time (recall that the fundamental and second harmonic have orthogonal polarizations), not to mention only 1 wavelength at a time, there will necessarily be some loss in one of the two wavelengths. On the other hand the faces can be AR coated for both wavelengths. Other workers in the field [Ber94] have determined that if the faces are Brewster cut for the fundamental, the small loss due to reflection of the second harmonic is more than made up for due to the larger circulating power in the fundamental. In other words, the small increase in the fundamental achieved by using the Brewster face leads to a large increase in circulating power due to the enhancement of the cavity. The second harmonic is even more dramatically effected since its power goes as the square of the fundamental power. If one plans to use the




Figure 4.4: The BBO crystal viewed from above. The angle θ_A is the cut angle. The angle θ_B indicates that the beam is incident on the crystal face at Brewster's angle. Brewster's angle for a beam that is to propagate along the crystal axis is achieved by cutting the crystal at $\theta_A = \theta_B$.

crystal single pass with no enhancement cavity, then AR coating for the fundamental and second harmonic is probably the best way to go. If the crystal is to be placed in a cavity, the Brewster faces should be used. Of course this is only true if maximum power is required. If this is not the case, the added difficulty of having to keep the Brewster-cut faces free of water contamination (BBO is sensitive to 40% humidity and above) may not be worth the effort.

The index of refraction for BBO at 462 nm (ordinary ray) is $n_o = 1.68214$ [Kat86]. If we take $n_{air}(462) = 1.00059$ [CRC66] then Brewster's angle at 462 nm is $\theta_B = 59.26^\circ$. Figure 4.4 shows a diagram of the Brewster cut BBO crystal from the top view. To obtain a refracted beam that travels parallel to longitudinal axis of the crystal, the angle θ_A in the drawing is equal to θ_B . Table 4.1 summarizes the pertinent parameters for the Brewster cut BBO crystal that will replace the AR coated crystal.

Table 4.1: The pertinent parameters for the Brewster cut BBO Crystal.

Indices of Refraction

$$n_o(461.2nm) = 1.68214$$

$$n_o(230.6nm) = 1.80898$$

$$n_e(461.2nm) = 1.56029$$

$$n_e(230.6nm) = 1.64614$$

Brewster's Angle

$$\tan\theta_B(462nm) = \frac{n_o(462)}{n_{air}(462)}$$

$$\theta_B = 59.26^\circ$$

Phase Matching Angle

$$\sin^2\theta_m = \frac{n_o^{-2}(\omega) - n_o^{-2}(2\omega)}{n_e^{-2}(2\omega) - n_o^{-2}(2\omega)}$$

$$\theta_m = 60.05^\circ$$

Walk-off Angle

$$\tan\rho = \frac{1}{2}n_o^2(\omega)(n_e^{-2}(2\omega) - n_o^{-2}(2\omega))\sin 2\theta_m$$

$$\rho = 90.52mrad$$

Cavity Waist

$$\omega_o = 20\mu m$$

Confocal Parameter

$$b = \frac{2\pi\omega_o^2}{\lambda_1} = 5.45 \times 10^{-3}$$

Crystal Length

$$l \approx 1.5 \times b = 8mm$$

Walk-off Parameter

$$B = \rho\sqrt{lk_1}/2 = 15$$

Aperture

$$a = 3mm \times 3mm$$

4.6.2 SHG Enhancement Cavity

Before going into detail on the SHG enhancement cavity (EC), I will give an overview of how it works and how it is locked to the fundamental beam. The schemes used have been described in the literature [Han80, Ber82] so the description here will not be comprehensive but rather should be viewed as complementary and focusing on our implementation. There are two components to the operation of the EC. First, is the enhancement of second harmonic and second is the locking of the EC to the fundamental so that the frequency can be tuned.

The enhancement of the second harmonic is achieved by making the circulating power at the fundamental inside the cavity greater than the power incident upon it. The amount of enhancement is determined by the losses in the cavity, so great care is taken to make the mirrors highly reflective. It turns out that the optimal input coupling is that which matches the losses in the rest of the cavity. Consequently, one usually obtains the best mirrors possible, calculates losses due to other intracavity elements such as that due to the doubling crystal and the output coupler, and then tries to match the input coupler to equal this value.

In our cavity there are two waists. The BBO crystal is placed at one of these which is just before the output coupler. The output coupler is chosen to be highly reflective in the fundamental and yet transmissive in the second harmonic. Consequently, the second harmonic does not circulate in the cavity and is not directly enhanced by it. Rather it is the fundamental that is enhanced. In this way the cavity need not have a high Q at two different frequencies. Once the fundamental makes a round trip, it now joins up with light entering the cavity. For these two to add effectively, they must be in phase and mode matched. The phase relationship is handled by tuning the length of the cavity. The mode matching is accomplished by placing two translatable microscope objectives at the entrance to the cavity.

It is simple to make a cavity resonant with incident light by changing its length. However any perturbation will knock the cavity out of resonance, not to mention the fact that the fundamental cannot be changed without losing resonance. Hence the cavity must be locked to the fundamental. This handles both noise shifts and intentional tuning. The locking scheme used, cleverly takes advantage of polarization sensitive elements in the cavity, and produces a discriminant with long tails to handle the inevitable frequency jumps and subsequent loss of lock.

The basic idea of the locking scheme is to translate a phase shift between the incident light and the light that would be resonant with the cavity into a voltage error signal. On resonance there is very little loss inside the cavity and so light passing through the cavity receives very little phase shift. However, off resonance the cavity has appreciable loss and there is an associated phase shift. So off resonance the light reflected off of the input coupler and the light which escapes out through the input coupler after circulating inside the cavity, are shifted in phase relative to each other. In addition, the polarization of each beam is slightly rotated with respect to the other due to the polarization sensitive nature of the BBO crystal. Consequently, when the reflected and transmitted beams are recombined the result is a beam with elliptical polarization. This is passed through a quarter wave plate which reduces the elliptical polarization into two orthogonal beams 90° out of phase and with different amplitudes. The orthogonal beams are then passed through a polarization sensitive beam splitter so that the beams can be detected separately in two photodiodes. The signals from the diodes are subtracted from each other producing the error signal. When the cavity is on resonance the quarter wave plate is rotated to produce equal amplitudes in the orthogonal beams and zero error signal. Any deviation from resonance will cause one amplitude to change relative to the other. The subtraction of the two signals gives a dispersive result. It is zero on resonance and opposite in sign above and below resonance. Hence, it serves as a good error signal and can be fed back to a PZT controlled mirror in the cavity to achieve lock. Perhaps the biggest advantage that this scheme has over other locking schemes is that the error signal produced has substantial "tails" that reach all the way to the next cavity fringe. Consequently, no matter how far the cavity gets knocked out of lock, it will always be able to reacquire lock automatically. In addition, the discriminant is created by the cavity as it stands—no additional apparatus (such as a device to frequency modulate the beam) is needed. All that is required is the $\lambda/4$ analyzer and the error signal detector.

4.6.3 Cavity Optics

The mode matching into the cavity is accomplished using two microscope objectives; one a $10\times$ and the other a $4\times$. These are placed on a mount usually used for the positioning of spatial filter optics, making it possible to finely tune the distance




Figure 4.5: Second Harmonic Generation Enhancement Cavity.

between them. The cavity itself uses a bowtie geometry (see Figure 4.5). For ultimate enhancement one must carefully select cavity optics for low loss. Initially however, our power requirements were not extreme and it was necessary to make do with optics already in the lab where possible. The input coupler (IC) is a flat dielectric coated laser mirror with 3.3% transmission (measured) at 462 nm. The second mirror is a high reflector (HR) and has better than 99% reflectivity. In our cavity the HR is positioned 171 mm from the IC. In between the IC and the HR is an EOM used to modulate the cavity on and off resonance. This EOM is not part of the enhancement or locking. It is used for the background suppression scheme that is described in Chapter 5 and will not be described further here.

The HR is tilted at an angle just large enough to cause the reflected beam to clear the EOM. This is the opening angle of the cavity. The larger the opening angle the more the beams are propagating off the axis of the curved optics which are in the cavity (to be described shortly). The greater the off axis propagation, the greater the astigmatism that is generated. Consequently, the smallest angle possible is desired. Since the Brewster faces of the BBO crystal introduce additional astigmatism, the

crystal can be oriented (plane of incidence equal to the cavity plane) so as to use the Brewster astigmatism to partially cancel the off axis astigmatism [Rau94]. This in turn constrains the input polarization to be in the plane of the cavity as well.

The next optical element is the small PZT mounted curved high reflector called the “tweeter”. The tweeter is placed 165 mm from the HR. We used a 699 tweeter provided to us by Coherent Corporation which has a radius of curvature of 75 mm and reflectivity better than 99%. The error signal generated by the cavity locking electronics is fed back to the PZT attached to this mirror.

After the tweeter is placed the output coupler (OC) which also has a 75 mm radius of curvature. Between the tweeter and OC, at the waist, is placed the BBO crystal. The OC is specially coated to be highly reflective at 462 nm (better than 99%) and highly transmissive at 232 nm (about 60%). Since the OC is transmissive in the UV, the second harmonic exits the cavity here and the cavity need only be resonant with the fundamental. The fundamental light reflected off of the OC is directed back to the IC where it joins incoming light. Some of the circulating light is leaked out of the IC collinear with light reflected off of the IC. Both of these beams, whose polarizations are not the same, are reflected into the polarization sensitive beam splitter and from there into the two photodiode detectors.

The total round trip length of the cavity is 57 cm giving it a free spectral range (FSR) of about 530 MHz. We have measured the cavity enhancement factor, e , using a technique suggested by Jim Bergquist [Ber94]. The value of e gives the ratio of the circulating power to the incident power. This means that the amount of UV generated by the cavity is e^2 times the power generated without it. Of course, losses at the OC and other optics outside must be considered as well. For example, we were able to achieve roughly 10 μW of UV single pass with 500 mW of incident power. An enhancement factor of $e = 15$ should give 225 times more power in the UV *with* the cavity or about 2.3 mW. This would be attenuated by about 50% due to the OC transmission and other output optics to about 1.2 mW. In fact, we have observed upwards of 1 mW, in basic agreement with this calculation. To measure e , a detector is placed after the HR and the power is measured leaking out of the cavity. The IC is then removed and the power is measured again. Let the first be P_{hr0} and the second be P_{hr1} . Then e is the ratio of the two:

$$e = \frac{P_{hr0}}{P_{hr1}} \quad (4.7)$$

This can be seen from the following. Let the incident power be P_i . From the definition of e , the power circulating in the cavity is eP_i . If T_{hr} is the transmission of the HR then the power leaking out of the HR will be $P_{hr0} = eP_iT_{hr}$. Now if the IC is removed, the cavity is no longer resonant and the power leaking out of the HR is simply, $P_{hr1} = P_iT_{hr}$. The ratio of P_{hr0} to P_{hr1} is now clearly e . We found e changed from day to day, no doubt dependent on how good our alignment was on that day. The best I was able to achieve with these cavity optics was $e = 24$. Others [Rau94, San93] have been able to get factors of 100 or more by employing better coatings on their optics.

As already mentioned, the BBO crystal is placed at the waist which is formed between the tweeter and the OC. The waist was numerically calculated to be $\omega_0 = 24.3\mu m$ with a confocal parameter of $z_0 = 5.5mm$. The waist was measured using a knife edge technique to be $\omega_{0h} = 17.9 \pm 2.5\mu m$ and $\omega_{0v} = 19.8 \pm 2.5\mu m$ where ω_{0h} is the waist in the cavity plane and ω_{0v} is the waist in the plane perpendicular. In this technique the distance required to change the beam attenuation (due to the knife edge) from 10% to 90% is measured. Call this value Δx_{10-90} . Using tables for the error function it is then possible to determine that for a Gaussian beam,

$$\omega_0 = \frac{\Delta x_{10-90}}{1.28}. \quad (4.8)$$

After leaving the cavity via the OC, the UV is passed through a prism to separate it from the fundamental traveling collinearly with it. This pure UV beam is now directed to the trap by a specially coated high reflector. Since the UV is the extraordinary ray in the crystal, it is polarized perpendicular to the fundamental. The fundamental emerges from the laser vertically polarized. It is convenient to configure the cavity such that the plane containing the mirrors is horizontal. Initially, when using the AR coated crystal, the fundamental polarization was perpendicular to the plane of the cavity (there were no Brewster faces to correct off axis astigmatism so there was no reason to rotate the fundamental polarization into the plane of the cavity) and the UV was perpendicular to this, or horizontal. The UV was p-polarized, or in the plane of incidence of the high reflector just mentioned, and the required angle was unfortunately quite near Brewster's angle. Careful coating of this optical element made high reflectivity still possible. In the subsequent version of the cavity which requires the fundamental to be in the plane of the cavity, the UV will now be s-

polarized with respect to the output reflecting optic, making the reflectivity less of a problem.

4.6.4 Cavity Locking Electronics

The circuit used to implement the polarization locking scheme is based on a schematic published in John Sandberg's Ph.D. thesis [San93] and also loosely on a schematic designed by Bloomfield [Blo83]. One of the cavity mirrors is mounted on a piezoelectric transducer (PZT) which is driven by the 3 pole filter contained in the circuit. Normally, a 3 pole rolloff above unity gain would not be stable, but as I will describe, Sandberg modifies the feedback to address this. The following is an overview of this circuit. See Appendix C.5 for the schematic.

Referring to this schematic, the input from the photodiodes is fed into an adjustable gain amplifier and then to an integrator network and a follower for monitoring purposes. Each of the integrators in the network can be reset by turning on the FET which is connected across the feedback capacitor. The first integrator in the network is straight forward. The second two, however, have a resistor in series with the feedback capacitor. At relatively high frequency the capacitor becomes a short and the integrator becomes a simple inverting amplifier. The purpose of this is to stabilize the circuit. At low frequency the three cascaded integrators give high gain. However, at high frequencies each integrator will contribute 90° of phase shift and if there were no resistors in the feedback the circuit would reach more than 180° of phase shift before unity gain. The resistors cancel the phase shift at high frequency because as inverting amplifiers these circuit elements now give exactly 180° of phase shift each or a net phase shift of 0° . This leaves just the single pole rolloff of the first integrator, thus acquiring only 90° phase shift and stable operation.

For the particular feedback values used, the cross over between integrator and inverting amplifier occurs at roughly 1 kHz. For frequencies below this the gain (in dB) is essentially the sum of the gains for the three integrators. For higher frequencies the integrators turned amplifiers now have a gain of 0 dB so the net gain is just that of the first integrator. Sandberg found that the gain was too high when the circuit was out of lock and it would just oscillate. To solve this he put in an RC network which disables the last two integrators for one time constant (about .5 s) after lock is lost. After this amount of time the circuit has settled down, and with small excursions,

high gain can now be applied without oscillations.

A switch determines whether the integrated signal or an internal generated ramp voltage is fed to the high voltage (HV) amplifier stage. The ramp voltage makes it possible to scan the cavity across resonance. This feature is essential for optical optimization. The switch is turned to scan mode and by adjusting the cavity optics one is able to achieve maximum fringe height and a TEM₀₀ mode. The output of the cavity is detected with a UV sensitive photodiode so that fringe height is a direct representation of power. Lowest order mode operation is initially obtained by reducing all but one mode per Free Spectral Range (FSR). Finally with the cavity in scan mode, the fringe is made as “smooth” as possible, getting rid of higher order modes which overlap the TEM₀₀ mode. The switch is then turned to lock and the integrated signal is fed to the HV stage. The op amp recommended by Sandberg is the Apex PA88. Its 450 V swing, 1 MHz bandwidth, and reasonable slew rate make it an ideal candidate for driving a PZT. However its price is relatively high at several hundred dollars. We used the much less expensive Apex PA41, which is a follow on to the PA88 that has 350 V swing, comparable bandwidth and better slew rate. This op amp is operated in single supply mode, but better frequency response can be obtained by using a bipolar supply.

As mentioned previously, the error signal has tails that extend out to the next (previous) cavity fringe. When the cavity goes out of lock it may go to a different fringe as it reacquires lock. Eventually, the integrators and amplifiers will no longer be able to provide gain for the potentially large error voltage. To remedy this, the circuit detects when the error signal is too large (either positive or negative) by running it through two comparators. Each of the comparators has a level that can be adjusted via an internal trim pot. If the error voltage is too large, a 555 timer chip fires a pulse which turns on the integrator feedback FETs, thereby resetting the integrators. Lock is then reacquired on the nearest fringe.

4.7 *Beam Shaping and Input Optics*

Having generated the UV light, it is now necessary to focus it into the trap. Since we do not have non-critical phase matching in BBO ($\theta_m \neq 90^\circ$) the UV beam that comes out of the cavity has a large degree of walk-off. In our case the walk-off is in the horizontal plane because that is the plane of the extraordinary ray. The result is that

outside of the cavity the beam is stretched into a highly non-Gaussian, rectangularly shaped profile. This beam can be difficult to focus tightly and as a result can create large amounts of scattering background as it passes through the trap. As has already been mentioned, in addition to walk-off there is astigmatism in the beam, due to the off axis propagation through the cavity (made worse in our case by the presence of the intracavity EOM which increases the cavity opening angles and places two Brewster windows in the path of the beam).

It is possible to greatly reduce the astigmatism. The first approach is to use Brewster faces on the BBO crystal in the enhancement cavity [Rau94]. To make this work the plane of incidence of the Brewster face must be the same as the cavity plane. In our case, we want the cavity plane to be horizontal so that necessitates rotating the polarization of the fundamental before it enters the cavity. Outside the cavity, standard techniques for correcting astigmatism can be used such as combining a spherical lens with a cylindrical lens. Astigmatism means that if a lens is placed in the beam, two orthogonal components of the beam will focus at different locations. If a cylindrical lens is also placed in the beam such that it is focusing one of the components and not the other, then adjustment of the various lens positions can make both components focus at the same point.

In contrast to astigmatism it is very difficult to correct the walk-off. An output lens is placed directly after the output coupler to minimize the divergence of the beam. The placement of this lens can also have a small effect on the aspect ratio (the ratio of the lengths of the long dimension to the short dimension of the beam profile) of the beam. One interesting approach that actually corrects the walk-off is to use two doubling crystals configured such that the walk-off in one is cancelled by that of the other [Bos89]. This technique is difficult to implement, especially in an enhancement cavity (not to mention that it requires two crystals) so we didn't pursue it. Techniques used for correcting astigmatism, such as the cylindrical lens approach or taking a normal lens and tilting it, seem to also improve the appearance of the unfocused beam shape. However, my experience was that these techniques did not have a dramatic effect on the shape of the focused beam and, more importantly, on the background. There was usually still a long and short dimension to the focused beam, even if the symmetry of the unfocused beam appeared better.

Another approach that we used was spatial filtering. We noticed that, to the eye,

the beam intensity seemed uniform across its profile, but the actual intensity was concentrated in one portion of the profile. We were actually able to get 60–70% of the light from 10% of the profile by using a partial beam block. Strictly speaking, this is not spatial filtering but did have a similar effect in that stray light is being removed. While this would improve our ultimate spot size (in both directions), it was very sensitive to beam alignment. Slight changes in beam position due to laser shifts or enhancement cavity shifts would have dramatic effects on how much light was being passed through the partial beam block. This necessitates constant vigilance in keeping the alignment correct and was deemed not worth the modest improvement. We also used “real” spatial filtering in which the beam is first tightly focused before passing through a pin hole. The hope is that the parts of the beam that are not tightly focused at the trap will also not be tightly focused at the pin hole and will therefore be blocked. While we found this to be true (there was a definite halo that was not being passed by the pin hole), it didn’t seem to have a dramatic effect on the background. Again, this approach was not made part of the apparatus because it didn’t seem worth the added complexity.

To determine improvement in beam characteristics, the ultimate test is how well the beam focuses through the trap. By placing a fluorescent card on the opposite side of the trap from the input beam we could observe it being cut off by the trap as we moved the beam transverse to its axis. Essentially, the trap serves as a knife edge and optimal focus occurs when the image is cut off with the shortest transverse beam displacement. In addition, at optimum focus, the direction of cutoff appears to rotate 90°. By adjusting the position of the enhancement cavity output lens relative to the output coupler to get a roughly symmetric output beam we were able to get a focused spot at the trap of roughly 20-40 μm in one direction and 60-80 μm in the other. While this was satisfactory, it is one area where significant improvement is needed.

The beam is directed from the SHG cavity to an expanding lens which is used to increase the spot size before the beam is sent into the final focusing lens. The focusing lens is placed at the base of the vacuum chamber and the beam is reflected up into this lens and focused into the trap. The beam is expanded before being focused so as to get the tightest possible spot at the trap. This means that a large part of the aperture of the focusing lens is used and it must be corrected for spherical and

chromatic aberrations. We already had such a corrected lens with an $f/5.6$ aperture from a previous experiment that was also transmissive in the UV. This aperture was sufficient to not limit our focused spot size. Since the focal length of the focusing lens is 85 mm and the distance from the lens to the trap is over 20 cm, the expanding lens also serves to adjust the focusing point to the location of the trap. During initial optical alignment, the position of the expanding lens is adjusted until the image of the beam is sharply cut off by the trap as it is moved transversely. Fine adjustments in the position of the beam, as it passes through the trap, are made by a two dimensional translation stage. The micrometers attached to the stage have micron precision and very little hysteresis. In addition to focusing the beam, the typical input alignment procedure consisted of translating the input beam with the 2-D stage. The amount of translation required to cut the beam off on one side, to cutting it off on the other. (in both directions) is used to center the beam in the trap.

4.8 Vacuum Chamber

The vacuum chamber containing the trap consists of a quartz tube with quartz windows fritted on (see Figure 4.6). It was manufactured by our glass blower Bob Morely. In fact, this was intended only as an initial design for testing purposes, however our final design using a stainless steel tube ended up being not as good. Only sapphire windows are available on a conflat flange for mounting on a stainless steel base. It is difficult (expensive) to get Sapphire windows welded to the stainless steel knife edge flange in sizes greater than about 1 inch. On the other hand, quartz windows in just about any size can be obtained and these are easily fritted onto a quartz tube. Large windows make large opening angles for output optics more accessible. The advantage of the stainless steel tube is that everything on it can be removed and replaced by simply unbolting it, whereas the quartz tube requires the glass blower to make changes. Nevertheless the need for large UV-transmissive windows was essential so this was the design used. In addition to the size issue, quartz windows have substantially better transmission than Sapphire at our wavelength.

Our tube consists of quartz plates fritted onto the main structure for the input/output windows and glass-to-metal seals attached via a knife edge flange for the pump and electronic feedthroughs. The input window is at the bottom and requires the beam to be directed upward to reach the trap. It is fritted on at roughly Brew-




Figure 4.6: Vacuum chamber.

ster's angle and the small amount of light that is reflected off of its surface is detected and used to monitor the background suppression modulation (see Chapter 5). The window directly above the trap is used for beam focusing and alignment using a fluorescent card. The windows in front and in back of the trap are the output windows. The plane of the trap is positioned so as to reflect light coming from the input window away from the primary output window which is in front of the trap. This tends to reduce the amount of background light scattered into the detection apparatus.

The electronic feedthrough consists of an off-the-shelf glass-to-metal seal which we then had modified by Morely to include 9 tungsten pins. On the vacuum side the pins are attached to the various trap parts. At the base of each pin is a mushroom shaped glass washer that protects the pins from being shorted to each other by deposits from the oven. On the air side, the wiring to control the trap is connected (see Section 4.9 for details). Great care is taken to avoid applying sideways pressure to the leads. Since they are made from Tungsten they are brittle, and in addition, the glass at the base of each one can easily crack. To cut them to the precise size that we want, we use an electric discharge machine (EDM) which cuts via a discharge between the object to be cut and a cutting tip. A servo loop keeps the discharge current constant so the cutting tip never actually touches the lead and no force is applied.

On the opposite side of the tube from the electronic feedthrough is the ion pump attached with another glass-to-metal seal. We used a Varian Triode Ion Pump which is optimized for Noble gasses. Carefully following good cleaning and bakeout procedures (see Section 4.8.1) we were routinely able to reach pressures in the 10^{-11} Torr range. The emphasis on Noble gas optimization was to accommodate our early work with buffer gas (He) cooled ion clouds. It was necessary to be able to effectively pump out the buffer gas to search for single laser cooled ions and this pump was able to go from 10^{-5} Torr of Helium to an overall pressure of 10^{-10} Torr within hours. In addition, our quartz windows are more permeable to Helium than glass windows. Consequently, even without the introduction of the Helium buffer gas, we still need the noble gas pumping ability. Due to the strong magnetic field generated by the pump, it was enclosed in an iron sleeve. This reduced the magnetic field at the trap from several Gauss to below the ambient Earth's magnetic field of roughly 1 Gauss.

To load buffer gas into the vacuum chamber, a Helium leak is connected to the side of the tube between the output windows and the pump. The leak consists of a

quartz finger attached to the side of the tube and contained in a quartz bulb. Since the permeability of quartz to Helium is very strongly dependent on temperature, a tungsten coil is wrapped around the quartz finger and current is run through it to let Helium into the tube. The coil is connected to a Variac via a step-down transformer supplying about 5 A of current. With this current the coil glows dull red. The bulb surrounding the coil contains about 100 Torr of Helium. We found that with 5 A running through the coil for about 10 minutes we were able to introduce about 10^{-5} Torr of Helium into the vacuum chamber. See Section 4.8.2 for more details on the Helium leak.

4.8.1 Cleaning and Bakeout Procedure

In order to reach the ultra-high vacuum required for single ion trapping it is necessary to thoroughly clean and bakeout the vacuum chamber. I will outline the procedure we use here.

The cleaning of the tube itself was left to our glass blower as part of the tube fabrication. The glass to metal feedthrough was placed in an ultrasonic cleaner, first in a solution of trichloroethylene to remove oils and other organic compounds, second in methanol to remove the trichloroethylene and lastly in acetone to remove the methanol and possible water deposits. Great care must be taken to make sure that no Acetone remains in small hard to reach places (such as inside the ovens). A heat gun is quite useful in getting rid of residual solvent. In general, new items such as the pump flange, are not cleaned since the manufacturer ships them ready for ultra-high vacuum. However, anything that has been touched or worked with should be cleaned using the same sequence of solvents. The exception, of course, is any part that has already been used at vacuum successfully and has not been contaminated.

After cleaning is complete, the pump was connected to the tube (using a moly-disulfide anti-seize lubricant on the screws to make later removal possible). A "T" is connected to the pump flange and the tube. An OFHC copper pinchoff and a Nude Bayard-Alpert vacuum gauge are all connected to the "T". As a rule of thumb, a knife edge flange should be tightened to the point where the two faces just touch. The screws should be tightened down in "cylinder head" fashion to keep pressure evenly distributed. If this is done correctly, then a relatively small torque will suffice to tighten the screws until mating occurs. At that point it will become significantly

harder to tighten any further and the job is done. If it becomes hard to tighten down before mating occurs than the copper gasket is probably not seated correctly.

When all parts are connected, the assembly is leak tested. The leak testing device is attached to the tube through the pinchoff (which is still open at this point) via a liquid nitrogen trap. All attachments are made with rubber hose and hose clamps. Very little, if any, vacuum grease is used on any of these connections. If it is used, great care must be taken when bringing the tube up to air after the test. If significant negative pressure exists when the rubber tubing is removed and grease is nearby, some can be sucked into the tube. The trap is necessary to prevent oil from the mechanical pump in the leak tester from contaminating the tube. Once the leak detector has been cycled up to its best vacuum (usually there are several stages) Helium gas is then blown over the tube. If no leaks are detected, then a needle is attached to the Helium hose, and Helium is blown directly into the leak test ports on all of the flanges. If no leaks are found, then the whole system is placed in a loosely fitting plastic bag which is then filled with Helium. If the system gets this far, then leak testing is done.

To remove the leak tester, the rubber hose between the liquid nitrogen trap and the tube is clamped. The cold trap is then removed from the nitrogen dewar and brought up to room temperature while still being pumped on. Once it is clamped, the hose connected to the tube can now be removed. This is done slowly so as not to create too much of an in rush of air.

After the system has been successfully leak tested, it is then connected to the bakeout system. Our system consists of a sorption pump as the fore-pump and a 30 l/s ion pump. The sorption pump is first cleaned by attaching a heating mantle to it and baking it for about an hour (a release valve on the sorption pump allows the resultant out-gassing to escape). After the mantle has been removed and the sorption pump has cooled somewhat, a liquid nitrogen dewar is attached to it and filled with liquid nitrogen until boil off subsides. This is all done with the sorption pump closed off from the rest of the system. The pinchoff on the tube is connected to the bakeout port and a valve to the sorption pump is opened bringing the tube to moderate vacuum (several mTorr). On the second and later tube recyclings this last step was sometimes hindered by the presence of Helium in our system. Since Helium is far from condensing at liquid nitrogen temperatures, we found it necessary to remove residual Helium with a mechanical pump working together with the sorption pump.

The sorption pump is allowed to pump on the system for roughly 30 minutes until the pressure comes down below 5 mTorr. At this point, starting the ion pump can be attempted. At this pressure, the ion pump will be working very hard creating several hundred mA through its electrodes. The pump is unstable at this point in the process because the current creates heat which causes a pressure increase. However, before appreciable heating has taken place some pumping action will occur. Therefore, the ion pump is switched on and the pressure is monitored carefully. At first it will go down, and then due to the heating it will begin to rise. As soon as the pressure starts to rise, the ion pump is turned off and allowed to cool. A fan, positioned between the sorption pump and the ion pump, blows air cooled by the presence of the liquid nitrogen onto the ion pump. This shortens the cooling time. This procedure is repeated several times until the pressure is low enough that the pump is stable (usually several mA of pump current). At this point the pump is switched to "run" and the pressure should come down to under $20\mu A$ of pump current within several hours (perhaps longer for a new tube).

The next step is to construct an oven surrounding the tube. We have several oven boxes whose enclosed volume we augment with fire brick. All cracks are "sealed" with fiber glass. Copper leads are attached to the filament feedthroughs so that a filament bakeout can also be accomplished in addition to the bakeout of the tube as a whole. The copper wire should have no insulation and be attached to the feedthroughs with alligator clips. A Chromel-Alumel thermocouple is also inserted into the oven, just touching the tube, to monitor temperature. The ovens are connected to a Variac which is set to 30 V to start the temperature increase. Both the temperature ($25^{\circ}C/mV$) and the pressure are monitored. The temperature should not rise faster than $1^{\circ}C/minute$ to prevent excessive thermal stress on the tube and the pressure should not go higher than 1-2 mA of pump current to keep the pump on. The Variac voltage is continually increased subject to these constraints until an optimal bakeout temperature of about $350^{\circ}C$ is reached. Done conservatively, this takes most of a day and may require an ultimate Variac voltage of over 100 V. Once the temperature has reached equilibrium at the desired value, the system is baked for several days or until the pump current comes down to roughly $20\mu A$. On one system which had never been put under vacuum this took over two weeks! Typically 3-5 days is sufficient, especially with carefully recycled systems.

When the system is finally cooled down, the same $1^\circ/\text{minute}$ change should be adhered to. As the ovens are cooling, the filament wires are connected to a Variac via a transformer able to supply 5 A. A small opening is made in the oven and the filament Variac voltage is turned up slowly until the filament glows yellow (roughly the temperature it will be when it is used to create an electron beam for ion loading). An inductive current meter is used to measure the current through the filament wires. Not more than 5 A of current should be required to reach the desired filament glow. Once this filament temperature is reached it is maintained for about 30 minutes.

With the filament bakeout complete and the ovens down to room temperature, the ovens are removed and a gauge controller is attached to the tube's pressure gauge. The pressure should be in the 10^{-8} Torr region or better. Next the pinchoff connecting the tube to the bakeout system is pinched off forming a cold weld at its tip. While this is being done, the pressure in the tube should be monitored for any sudden changes—a sure sign that the pinchoff has been done incorrectly. The tube is now free standing and the power supply for its own pump should be connected and turned on. If everything has gone well the pressure should go to the 10^{-9} Torr or even the 10^{-10} Torr quickly. However, I have had systems that started at 10^{-8} Torr and eventually ended up at 10^{-10} Torr or better, but took several days of pumping to get there. At this point the tube is ready to be set up in the experiment.

4.8.2 Buffer Gas Plumbing

Initially, the Helium leak was sealed off once it was filled with Helium. This worked quite well as far as getting Helium into the system, but quartz retains some permeability to Helium at room temperature and so it continues to leak in even when no current is flowing through the coil. The room temperature permeability of quartz to Helium is:

$$G \approx 10^{-11} \frac{\text{cm}^3(\text{STP})\text{mm}}{\text{cm}^2/\text{sec}/\text{Torr}}. \quad (4.9)$$

In words, 10^{-11} cm^3 of Helium at STP flowing per second through a 1 cm^2 area 1 mm thick and with 1 Torr pressure drop between the two sides. The leak rate is then,

$$R_l = \frac{GPA}{t} \quad (4.10)$$

where A is the cross sectional area of the membrane separating the two regions and P is the difference in pressure between them. The quartz finger has a surface area of




Figure 4.7: Apparatus used for loading helium buffer gas into the vacuum chamber.

about 30 cm^2 . The thickness of the membrane is 1.5 mm and the pressure in the bulb is 120 Torr. Putting these numbers into eq. 4.10 gives, $R_l = 1.9 \times 10^{-8} \text{ Torr} - l/\text{sec}$. For our bulb which had a volume of about 320 cm^3 , this gives an expected change in pressure of about $5 \times 10^{-5} \text{ Torr}$ over one day. We have observed on the order of 10^{-4} Torr per day in rough agreement with this calculation. Obviously this is not a problem during times when we want Helium to be in the system, but when we want ultra-high vacuum it is disastrous.

We resolved this unwanted leakage by placing input and output ports on the leak bulb allowing us to load the bulb with Helium when we wanted it in the vacuum tube and removing it from the bulb at other times (see Figure 4.7).

The following procedure was used to operate this apparatus. First, it is necessary to evacuate the bulb so that when we put Helium into it and we measure the pressure, we will know that the pressure is entirely Helium. With all valves closed the pump is turned on and the cold trap is placed in its liquid nitrogen bath. Pump valve

2 is opened to evacuate the cold trap and then pump valve 1 is opened to pump out the bulb. With the pump operating, the Variac is slowly brought up to a high enough voltage so that the coil just begins to show dull red (with the 6 V step-down transformer in use, we set the Variac at 70-80 V). Care must be taken here because as the bulb is evacuated the coil is less able to dissipate heat and gets hotter. Sometimes the Variac must be turned down from its initial value to prevent the coil from overheating and causing damage. The heat provided by the coil allows the pump to remove Helium impregnated in the walls of the quartz finger. This bakeout lasts about 10 minutes.

Next, the bulb needs to be filled with Helium. The Variac is turned off and the coil is allowed to cool down for about 10 minutes. Pump valve 1 is closed to remove the pump from the system. The Helium tank regulator is turned on providing positive Helium pressure at the Helium valve which is then opened. Helium then floods the system up to pump valve 1 and the gauge can be used to estimate how much Helium has been allowed in. Initially we put only enough in to just make the gauge start to move (starting at about 50 mTorr). After much use, however, we found that allowing a significant portion of atmospheric pressure of Helium worked better and wasn't so much that it flooded the vacuum system. When the desired amount of Helium is in the bulb, the Helium valve is closed and the Helium regulator is turned off.

To put Helium into the vacuum tube, the coil is again turned on to the same voltage as before. We found that leaving the coil on for 5-10 minutes gave us about $10^{-5} - 10^{-4}$ Torr of Helium in the vacuum tube. At this point, pump valve 1 is again opened to remove the Helium from the bulb. The coil is left on as before, to get the residual Helium out of the finger. We generally continue to pump on the heated coil for up to an hour (since the Helium is now in the vacuum tube, other work can continue in parallel). Finally the coil is turned off and the pump continues to operate until the whole system is at room temperature.

4.9 *Trapping Apparatus*

During the course of the experiment several trap headers were built. The header consists of the glass-to-metal flange, the tungsten feedthroughs and the apparatus attached to them. Here I will restrict myself to a description of the two successful systems we used; one for trapping ion clouds and the other for trapping single ions.




Figure 4.8: Paul trap used to trap ion clouds.

I will have more to say on the other trap headers that we used in Chapter 6.

The header used to trap ion clouds consisted of a Paul trap (a toroidal ring and two spherically tipped endcap electrodes)—See Figure 4.8. The trap geometry is the same as that used by Gary Janik in his thesis on trapping Magnesium and Barium ions [Jan84, p. 43]. The ring electrode was constructed out of .035 in stainless steel wire bent around to form the ring at the end of a lead. The endcaps were made out of 1/8" stainless steel welding rod with the roughly hemispherical shape being achieved on the belt sander. The ring diameter, $2r_0$, and the endcap separation, $2z_0$, were both 1.61 mm. The endcap and ring leads were wrapped in, and spot welded to an advance strap. The advance straps were spot welded to the tungsten feedthroughs. An RF drive with frequency $\Omega = 2\pi \times 10\text{MHz}$ and amplitude $V_0 \approx 400\text{V}$ is applied to the ring electrode to give a secular frequency of $\omega_0 = 2\pi \times 1.8\text{MHz}$ and a well depth of $D_0 = 26\text{eV}$.

As has already been mentioned, each feedthrough is fitted with a small glass “mushroom” to prevent shorting. The mushroom is held in place by a small piece of

.01 in Tantalum spot welded to the tungsten lead.

The ovens consist of .001" stainless steel shim stock rolled into a cylinder by wrapping it around an extra piece of tungsten rod (the same used for the feedthrough which is 12 gauge). The shim stock is first spot welded to the rod, then wrapped around the rod and then spot welded where it overlaps itself. By gently twisting the rod, the inner weld is broken without breaking the weld of the shim to itself and the oven cylinder is removed from the rod. The oven is about 1/4" long and about 1/16" diameter. A small portion of one end is crimped and spot welded to a .03" Tantalum lead. Another piece of shim is wrapped around the open end of the oven and spot welded to a second .03" Tantalum lead. A .01" Tantalum "hatch" is attached to the open end of the oven and bent across the opening after the Indium is loaded to prevent it from falling out. The two Tantalum leads are then spot welded to tungsten feedthroughs via a 1 mil advance wrap (tungsten and Tantalum don't spot weld to each other very well but each spot welds to advance easily). About 40 J of energy is required to spot weld 12 gauge tungsten to advance and 30 J to spot weld .03" Tantalum to advance. A small strip of Indium, which comes in wire form with much the consistency of solder, is inserted into the oven cylinder (after the whole apparatus has been cleaned) and the Tantalum hatch is closed. Due to the high vapor pressure of Indium a current of only 4 A run through the oven was usually enough to load ions. Note that the melting point of Indium is about 150°C so it is important to configure the oven in the vacuum chamber such that it never has its opening pointing downwards. This is true for the configuration on the optical table *and* on the bakeout system.

The current running through the thin filaments generates electrons which are used as an ionizing electron beam source. Two .03" Tantalum wires were connected to Tungsten feedthroughs in the usual way and bent around so that their ends were parallel to the feedthroughs and about 1/4". apart. The filament itself is made from .008" 2% Thoria-impregnated Tungsten (to enhance emissivity). It is first etched in a solution of Potassium Hydroxide to get it's diameter down to about .006". A section of this is cut to fit between the two Tantalum holders and is carefully spot welded to them with about 14 J. It is very easy to break the filaments with too much welding energy. Since resistance is highly temperature dependent in these filaments, it is easy to burn them out by simply applying a fixed voltage across them immediately. A

constant current supply was built to provide an adjustable current of approximately 6 A. This resulted in about $25\mu\text{A}$ of electron beam current at the trap—plenty to ionize the Indium.

To direct the electrons being emitted from the filament, an anode with a small hole is placed between the filament and the trap. There is a direct line of sight from the filament, through the hole, to the trap center. The anode is kept positive relative to the filament, thus attracting electrons which are then directed towards the trap. In addition, the anode serves to protect the filaments from the ovens and blocks light given off by the filaments from getting into the detection optics. The anode was constructed out of a 1 cm square piece of 1 mil advance spot welded to a Tantalum lead and positioned so that it partially surrounds the filament. Usually the anode as well as the ovens and filament, are positioned on longer feedthrough pins than the trap so as not to block any of the laser or fluorescence paths.

Several modifications to the header were made for trapping single ions. We used a Paul-Straubel trap with no endcaps [Yu91]. The trap is formed from a single piece of $125\ \mu\text{m}$ (.005") Tantalum wire. It is first bent around a mandrill and then twisted around itself tightly. This forms the first lead. Next, a second mandrill of diameter equal to the desired ring diameter is placed tightly against the twisted wire which is then wrapped again to form the loop of the trap. The wire is then twisted around itself again to form the second lead. Thin wire was used so that the trap could be heatable, making it possible to drive off micromotion-causing impurities. The inner diameter of our trap was 1 mm and with $125\ \mu\text{m}$ wire the effective lead thickness was $250\ \mu\text{m}$.

A Paul-Straubel trap is only able to approximate the quadrupole potential provided by a Paul trap. The closer the ion is to the center of the trap, the better the approximation. The simpler design makes laser alignment easier, reduces scattering background, and makes trap heating possible. The price of these advantages is an increase in the RF drive amplitude required to produce the same well depth and a large anharmonicity at a given distance from the trap center. An exhaustive numerical analysis of these effects has been carried out on several trap geometries [Yu95]. In Section 2.2 I discussed the results of Yu's analysis which apply to a square ring trap. While this reasonably approximates the trap ring that we use, it doesn't take into account the leads connecting the ring to the feedthroughs. We chose to use a

Table 4.2: Parameters of the single ion trap. R_0 is the trap inner diameter, t is the thickness of the trap ring, d is the lead thickness, V_0 is the amplitude of the RF drive, Ω is the drive frequency, ω is the secular frequency, D_z is the well depth in the z direction (trap axis), β is the calculated modulation index for a DC potential of 1 V applied across the ring, L is the voltage loss factor for a Paul-Straubel trap compared to a Paul trap, and C_{z4} and C_{z6} are the anharmonic coefficients. As a side note, V_0 is measured by disconnecting any ground connections to the trap, turning on the filament as an electron source and turning on the trap drive. Charge builds up on the trap in accordance with the drive amplitude and the latter can be measured with a high impedance probe attached to the trap electrode.

$$\begin{aligned}
 R_0 &= 1mm \\
 t &= .005in. \\
 d &= .010in. \\
 V_0 &= 500V \\
 \Omega &= 2\pi \times 9.95MHz \\
 \omega &= 2\pi \times 1.04MHz \\
 D_z &= 3.3eV \\
 \beta &= 68 \\
 L &= 5.37 \\
 C_{z4} &= .105 \\
 C_{z6} &= -.687
 \end{aligned}$$

trap that had a relatively good compromise between voltage loss and anharmonicity in the analysis of the ring alone as a starting point. We then calculated the effect of adding the leads to this geometry. The geometry we chose uses 5 mil (125 μm wire) forming a 1 mm diameter ring. Since the leads are twisted wire the effective thickness is 250 μm . With a radius of 500 μm this gives an aspect ratio of $A = t/r_0 = .5$ where t is the thickness of the wire in the ring. When the analysis is carried out on this geometry (including the leads), we get [Yu95a], for the fourth order term, $C_4 = 0.105$, for the sixth order term, $C_6 = -0.687$, and for the loss factor, $L = 5.37$. These are large anharmonicities but the loss factor is reasonably good. See Table 4.2 for a summary of trap parameters.

The trap is mounted to the tungsten feedthrough electrodes at an angle to break the radial symmetry enabling the laser to cool in all three dimensions at once. See




Figure 4.9: The Paul-Straubel trap used for single ions. Also shown is the mounting angle used to break the trap symmetry. Note that the trap leads don't show the wire being twisted—difficult to do in my drawing program!

Figure 4.9. It is also coated with Aerodag, a flat black carbon layer which, it was hoped, would reduce background.

In addition to the filament, field emission points (FEPs) were added to the single ion trap header. FEPs can be easily manufactured, produce reasonable electron beam currents and most importantly, produce *no* light background. We thought that using FEPs for our electron beam source would make it possible to load ions and search for an ion signal simultaneously since the weak signal would not be lost in the noise associated with the filament background. As it turned out, we almost always use the filament but I will give the FEP fabrication procedure here since it may be useful in the future.

The FEPs are made from 3/4 in sections of 20 mil *non-welding* tungsten. Place a tungsten section in a pin vise attached to a micrometer and set up the micrometer so that it can be used to lower the tungsten section into an etching solution. The usual solution is 8% KOH placed in a small beaker along with a metal plate electrode. A voltage is applied between the immersed tungsten and the plate electrode with a Variac that attached through a 6 V step-down transformer. A DVM is attached in series with the Variac so that the current can be monitored. Insert the tungsten into the solution .3 in beyond where the meniscus forms and raise the Variac voltage until 360 mA of current flows. The tungsten will preferentially etch at the sharper edges naturally producing a sharp point. The amount of current is proportional to the surface area so as etching proceeds, the current will go down. When it has reached 180 mA (about 10 minutes) immediately turn the voltage down until only 30 mA of current flows. Allow the solution to continue etching until the current is further reduced to 20 mA and then turn the Variac off. Remove the tungsten from the solution and wash with distilled water. View the tip under a microscope. If the etching has been done correctly, the tip should disappear into the micron range. Such a tip can produce greater than 1 μA of electron beam current when a 500 – 700 V potential is applied. I attached several FEPs to the trap header by spot welding them to a .01 in Tantalum lead (a tricky business at best since the slightest touch of a FEP tip will destroy it) to form a sort of “Gatling gun” arrangement. For any given voltage applied to the FEP electrode, only one FEP will fire (due to the highly non-linear turn on characteristics). So, including others is simply insurance against one burning out.

The ovens and anode are the same as for the ion cloud header except that the anode is now used to direct a beam from the FEP as well as the filament and so two holes appear in it. Since the trap has no endcaps, small compensation electrodes are added. These consist of small 5 mm square pieces of advance spot welded to .01 in Tantalum wires which are welded to the feedthroughs. The compensation electrodes are placed perpendicular to the axis of the trap such that their centers are “pierced” by it. They are far enough away that they don’t interfere with the open Paul-Straubel geometry but close enough to shift the trap center when as little as several volts are applied (about 3-5 mm from the trap center).




Figure 4.10: Block diagram of electronics used to generate the trapping potential.

4.9.1 Trapping Electronics

The electronics used to generate the trapping potential are shown in Figure 4.10. The RF drive is generated by an HP 606A signal generator producing about 2 V at 9.95 MHz. Even though these generators are old, they are quite stable, and they have been used in this lab for this purpose with good results. The amplifier is built around an IRF 510 VFET and gives about 20 dB of gain. The schematic can be found in appendix C.1. Following the amplifier the RF is passed through a standing wave ratio meter (SWR) which gives a reading of how well matched the various components in the circuit are by comparing the forward power to the reflected power [ARRL, p. 16-2].

$$SWR = \frac{1 + \rho}{1 - \rho} \quad (4.11)$$

where $\rho = \sqrt{P_r/P_f}$, P_r is the reflected power and P_f is the forward power (the lower the SWR the better). The RF is matched into a step-up coil and the frequency is fine tuned to optimize the SWR. We usually operate at a value of $SWR = 1.5$. The output of the step-up coil is coupled directly to the trap and capacitively divided into an HV monitor. The step-up coil provides another 20-30 dB of gain for an amplitude of approximately 400-500 V at the trap. The HV monitor value has been calibrated using the technique already mentioned of charging up the trap. The HV amplitude at the trap is 149 times the amplitude measured at HV monitor.

The oven is very simply controlled by a Variac. Of course the actual current flowing changes from oven to oven, but we found roughly 4 A sufficient for most of our needs. See Chapter 6 for more details on oven operation and the loading process. The filament is operated by a constant current power supply. Since the resistance of the filament changes with temperature, if it's power supply is simply set

to a predetermined voltage with no constraints on the current, the filament can be easily burned out. The constant current supply automates the process of bringing the filament up to the correct voltage without causing damage. To turn on the electron beam, a second switch is supplied which lowers the potential of the filament relative to the anode by 100 V. Thus, electrons emitted by the filament are now attracted to the anode, and go through the anode hole which directs them towards the trap. When the electron beam is turned off, the potential of the filament is simply raised 30 V above the anode. Since the filament takes about 1 minute to warm up, the electron beam switch allows the operator to turn the electron beam on and off at will without having to wait for the warm-up. A microammeter is placed in series with the electron beam switch and the filament making it possible to measure the electron beam current (of course not all of this current goes directly to the center of the trap). The compensation plates and the anode were generally kept at ground.

Much of this apparatus was placed under control of the data collection computer while retaining the option of manual control. A Forth program called SCAN was written to automate ion loading and spectra taking. When SCAN is started, it turns on the oven and then waits for the number of seconds indicated in the command-settable variable, OV-WAIT. After this time expires, the filament is turned on and again the program waits, now for the time indicated in the variable FIL-WAIT. After this time, the electron beam is turned on and then after the time indicated in EBEAM-WAIT, everything is turned off. At this point, the loading process is complete. SCAN then starts the laser frequency sweep and displays the detected counts on the data collection computer screen. After a full screen of data has been collected. SCAN saves the associated file and terminates. This program makes it possible to perform the entire loading process, which might take 10 minutes, by entering just one command. Other useful programs include LOAD-ION, which just does the ion loading part of SCAN; DXA, which can be used to display a screen of data with blocks of channels averaged; and ADDSCR, which adds two screens together—a manual integration. These last two programs were written largely to assist in detecting small signals.

4.10 Detection Strategy

Detection of the ion signal was done in several stages. First, a buffer gas cooled ion cloud was detected electronically. The coherent oscillations of the cloud in response

to a drive voltage causes a small current to be induced in a tank circuit. Next, a buffer gas cooled ion cloud was detected optically by exciting it with the laser and collecting the resultant fluorescence in a photomultiplier tube. Finally there was optical detection of a single ion. In what follows I will describe the apparatus used for each of these cases.

4.10.1 *Electrical Detection*

Basically, the idea in electrical detection is to apply a ramp voltage to the ring in addition to the RF already being applied. A second RF signal is applied to the endcaps to drive ion secular motion. The ramp causes the secular frequency to change and as it comes into resonance with the endcap RF, a large coherent cloud oscillation is created which induces a current in an LC tank circuit attached across the endcaps. This current is picked up and amplified to produce the electrical detection signal. This type of detection was particularly useful in the early stages of the experiment to demonstrate that our loading mechanism worked (before the laser system was even set up).

Figure 4.11 shows a block diagram of the apparatus used for electrical detection. The normal RF drive at Ω applied to the ring of the trap is augmented by a ramp voltage at several hundred Hz. With just Ω present the secular frequency of the trap would be ω_0 . The addition of the ramp shifts the secular frequency to $\omega_0 + \epsilon$. Meanwhile a second RF drive at $\Omega + \omega$, where $\omega \approx \omega_0$, is applied to the endcaps. When the secular frequency of the trap, $\omega_0 + \epsilon$, is equal to ω , the entire cloud will perform a driven oscillation and its amplitude will grow until it reaches equilibrium with various damping mechanisms present. The large amplitude coherent oscillation at ω of the ion cloud will induce a current in the LC tank circuit attached to the endcaps. This signal at ω is fed into a preamp and then into a phase detector. A schematic for the preamp shown in Figure 4.11 can be found in appendix C.2. The phase detector derives its quadrature information from an ω signal obtained by mixing the original Ω ring drive with the $\Omega + \omega$ endcap drive. Finally, the phase detector produces a DC signal proportional to the amplitude of ion oscillations at ω which is further amplified and displayed on an oscilloscope (the scope is configured in XY mode with the ramp voltage being used as X and the signal as Y). The result is a roughly Lorentzian line shape and is a good indicator of whether there are ions in




Figure 4.11: Block diagram of electronics used for electrical detection of an ion cloud.




Figure 4.12: Apparatus used for optical detection of ions.

the trap.

The anharmonicity of the trap can be demonstrated by making the well depth smaller. This causes the amplitude of the ion oscillations to increase, thereby moving the ion cloud further from the center of the trap where it does not have a perfect quadrupole potential. The result is a frequency shift and an increased asymmetry in the electrically detected lineshape [Lan76, p. 87]. The asymmetry is readily observable on the oscilloscope.

4.10.2 Optical Detection

Optical detection of an ion cloud is accomplished by exciting the cloud with the cooling laser focused into the trap and resonant with the cooling transition at 231 nm. The ion fluorescence is collected by a spherical mirror which focuses it into a photomultiplier tube. Figure 4.12 shows the apparatus used. The ion fluorescence emerges from the ion cloud in a dipole pattern with an axis dependent on how the polarization of the exciting beam is oriented. In our case the dipole axis is horizontal and perpendicular to a line connecting the trap with the output mirror. A small

portion of it goes through the opening between the trap ring and trap endcap. This light forms a cone which is collected by a 6 in diameter $f/1$ spherical mirror. The mirror is positioned at roughly its radius of curvature (12 in) from the ion source so that light from the ions will be focused onto a photomultiplier tube placed just below and in front of the trap (outside the vacuum). This makes the collection efficiency of the mirror effectively $f/2$ and therefore the collected output cone forms a solid angle of $1/64 \times 4\pi$ steradians ($1/64$ of the sphere).

Ideally, a fast corrected lens should be used on the output. These can be quite expensive, especially if they are to be UV transmissive. As an intermediate step, we chose to use a mirror instead. The mirror is about 2 orders of magnitude less expensive than an equivalent aperture corrected UV lens, however it does have disadvantages. First, our mirror is Aluminum with a Magnesium Fluoride coating. This is inexpensive but has a reflectance at our wavelength of about 85% degrading somewhat over time. This is certainly no limitation of the mirror approach, since a dielectrically coated (more expensive) mirror could be used to get reflectance up into the 99% range. The next concern is that the collected light is off axis. That is, the light is not reflected directly back towards the trap, but slightly off to one side. This can introduce astigmatism and coma into the reflected beam. Finally the distance from the PMT to the mirror is not the same as the distance from the mirror to the trap causing additional aberrations. Nevertheless, this approach had been used successfully in this lab and so we decided to try it.

Just before entering the PMT, the fluorescence passes through an "astigmatism correction plate". The plate is simply a quartz window tilted at an angle with respect to the axis of the fluorescence. The tilt angle introduces additional astigmatism which can be used to cancel that caused by the off axis reflection. The correct angle is determined with the following procedure. Fluorescent paint is placed on the side of the pinhole plate covering the PMT aperture. The UV laser beam being focused into the trap is tilted off the center of the trap until it is reflected off of the side of one of the trap electrodes. The reflection now serves as a sample UV point source from very near the trap center and can be used to get the output fluorescence focused into the PMT. Since the reflected light is off axis, the induced astigmatism produces two different focusing points in two orthogonal directions. The spot on the fluorescent paint never focuses to a sharp point but goes from diffuse to elongated in one direction, to diffuse

and then elongated in the orthogonal direction as the PMT is translated along the light axis. The closer in space the two elongation points are, the less astigmatism there is and the sharper will be the focus in between them. If the astigmatism is bad then not only will some fluorescent light be lost, but light being reflected from the trap may get in, thereby increasing instrumental scattering background. The angle of the astigmatism correction plate is determined empirically by finding the point where the two orthogonal foci are closest to each other. We have been able to use this system to reduce a spot of several hundred microns to under 50 microns. While this is a nice improvement, we still have a large background signal and it is most likely caused by instrumental scattering getting into the PMT via coma (which is not removed by the correction plate) and residual astigmatism.

After passing through the correction plate, the fluorescent light is then collected by a PMT detector. The detector we use is a Hamamatsu R1657 side-on type photomultiplier tube. The photocathode material is Cs-Te which the manufacturer calls solar blind. While we did find the detection to be greatly attenuated in the visible, it is not completely cut off. It was still quite necessary to run with the room lights off and our PMT shielded from ambient light sources. It is possible that our sensitivity to room light was really to UV which is present in room light (particularly fluorescent lights)—we never pursued this enough to find out. The quantum efficiency for this PMT at 231 nm is roughly 30%. Once the signal is collected by the PMT it is amplified, discriminated and then sent to a pulse generator to create a TTL pulse for each photon count detected. This TTL signal is passed along to the computer for display and storage.

A pinhole is placed over the PMT aperture to prevent scattered light from being detected along with the signal. We used pinhole diameters between 50 μm and 200 μm . As has already been mentioned, we put fluorescent paint on part of the pinhole surface for alignment purposes. The presence of the correction plate complicates alignment. Typically, alignment is accomplished by back-lighting the trap and centering the reflected shadow of the trap on the pinhole. Normally, the fact that the back-light source and the actual fluorescence are at very different wavelengths wouldn't make a difference. However, the correction plate is tilted and so it refracts the beam (exactly the feature we are taking advantage of for correction). Of course the fluorescent beam is refracted differently than the back light so it is not correct

to center the shadow on the pinhole. Instead we calculated that the difference in transverse position of the two beams due to refraction is about $400 \mu m$ (see Appendix B.1). We were able to determine this empirically by observing the position of the optimal ion cloud signal and it agreed well with our calculation. More details on the alignment procedure will be covered in Chapter 6.

In order to calculate an expected ion signal it is necessary to track all of the various loss factors present between the generation of ion fluorescence and the detection apparatus. Combining all of these we get an overall collection efficiency for ion clouds:

$$E_{cl} = \Omega E_{pmt} T_w R_m A_{ph} T_{cor} A_{int}. \quad (4.12)$$

The quantity Ω is the ratio of the detection solid angle to the sphere (with the approximation that radiation is isotropic instead of in dipole pattern—for our purposes here this is a reasonably good approximation) and has a value of $1/64$. E_{pmt} is the quantum efficiency of the PMT and is taken to be 30%. T_w is the transmission of the 4 quartz window surfaces: two for the vacuum chamber output window and two for the PMT window. We used $T_w = (.96)^4$. The quantity, R_m , is the reflectivity of the output mirror and was taken to be .85. The quantities A_{ph} and A_{int} taken together represent what portion of the fluorescing region of the ion cloud is being observed. A_{int} is the fraction of the cloud that is interacting with the focused laser beam. Assuming a Gaussian beam (an approximation) with a $20 \mu m$ waist intersecting a spherically shaped ion cloud of diameter $z_0/\sqrt{10}$, we calculate $A_{int} = .00235$. A_{ph} is that fraction of A_{int} which is not cut off by the pinhole and for a $75 \mu m$ pinhole we get $A_{ph} = 1/6$. Finally, T_{cor} is the transmission of the correction plate and is about .7. All of this taken together gives:

$$E_{cl} = 9.2 \times 10^{-7}. \quad (4.13)$$

For single ions, A_{int} and A_{ph} do not enter in since we are dealing effectively with a point source. Consequently in this case we have:

$$E_s = \Omega E_{pmt} T_w R_m T_{cor} = 3.9 \times 10^{-4}. \quad (4.14)$$

I have already mentioned that the photon counting data is sent to a computer for display and storage. The computer used is a very simple Apple IIe. The decision

to use this computer was based on several factors. First it was possible to obtain working Ile's from our University surplus for \$25. At this price we purchased several and had no shortage of spare parts. Second, we have an ongoing need to be able to modify whatever data collection system we are using at a hardware level to interface to various sources of data. Extensive experience exists within the lab making these sorts of modifications to the Ile and so it was relatively easy to simply copy what had already been done. The primary modification made in our case was the addition of an interface board which reads the TTL photon counting signal and determines the total number of counts detected per some predetermined unit of time. This information is placed in registers from whence it can either be stored or displayed on the screen as a function of time. In addition the board generates a ramp voltage which is sent to the laser external control for frequency scanning. For each unit of time (a so-called "bin" whose value is set by the operator), the voltage is incremented by $1/256$ of the full scale value to give 256 bins for the whole screen. Thus, when the computer is causing the laser to scan its frequency linearly in time, the photon counts vs. time display on the computer becomes a photon counts vs. frequency display, or spectrum. Far more sophisticated data acquisition systems exist, but none cheaper or as simple to set up. The automatic loading programs described in Section 4.9.1 are also on the Ile. Data was easily transported by serial port to an Apple Macintosh system for analysis.

4.11 Background

The subject of background is a thorny one for this experiment. The elimination of one of the sources of background required some hard work and a little bit of cleverness and was a key to our success. Essentially there are two sources of background. First, there is ambient light which comes from room light or any other surrounding light source. It is usually quite simple (though perhaps not convenient) to eliminate this type of background by enclosing the trap and detection apparatus in a box which seals out light, and by taking advantage of the fact that the fluorescent light is at a different wavelength than the ambient light (e.g., the use of solar blind detectors). The second source of background is caused by scattering of laser light off of parts of the trap and is called instrumental scattering. This light is at the same wavelength as the fluorescence and comes from almost the same location so that the techniques used to eliminate ambient background won't work to reduce the scattered variety.

Technically, scattered light can get into the detection system from other parts of the experiment as well, but these are usually far away from the trap and can be shielded by the same apparatus used to screen out ambient background. So when dealing with instrumental scattering background we are talking about light scattered from a region very close to the ion source—probably the trap electrodes.

Ideally, the pinhole in front of the PMT aperture should screen out even instrumental background and in fact it does a fairly good job of doing just this. For example, if the output optics are set up to give 1:1 magnification, then the image of the trap and the ion source at the PMT will be the same size as it is at the source. If the pinhole is made smaller than the inner diameter of the trap and the optical imaging is perfect (no aberrations) then, by definition, no light except that coming directly from the ion can get into the PMT. However no imaging optics are perfect and this is especially true in our case where we have aberrations due to the geometry of our optics. In non-perfect optical systems, points are not imaged as points. Rather, they are imaged as a spread-out spot. These spots may be small and circular in the case of diffraction limitations, or they may be large and irregularly shaped in the case of astigmatism and coma. In this latter case, one point on the trap electrode and one point in the center of the ion cloud might be imaged to two spots which overlap. The imaged spot belonging to the point on the trap will be centered outside of the pin hole, but because of its large irregular shape, some part of it is well within the pin hole. Hence, we have introduced a rather insidious form of background at the same wavelength as the fluorescence and we have no simple way of blocking it.

It turns out that with In^+ even the instrumental scattering background can be virtually eliminated (at the price of a relatively small amount of signal). This is the subject of the next chapter on background suppression. Suffice it to say here, that by sealing the trap from ambient light and by suppressing scattered background, we were able to reduce our background signal from tens of thousands per second to under 30 counts per second. This was essential to seeing a signal which was initially only 50 counts per second.

Chapter 5

BACKGROUND SUPPRESSION

In the previous chapter I pointed out that background from light scattered off of the trap electrodes needs to be reduced as much as possible. If one is detecting a strong signal, background can often be subtracted out leaving the signal largely untouched by excessive noise. However, if one is detecting a weak signal then the large amount of noise that will accompany a large background signal may make it difficult to analyze or even detect the signal. Since background noise usually goes as the square root of the background amplitude (shot noise), when the background is high enough that the signal to noise ratio (SNR) is 1, then the signal will become increasingly difficult to pick out. This turned out to be the case for the In^+ cooling transition. We needed a way to reduce background, or more importantly, increase the SNR, by a significant amount.

One approach is to reduce the amount of scattered light that reaches the detector without significantly reducing the signal. This can be achieved with a judicious choice of PMT pinhole size. We eventually found that for single ions it was possible to make significant improvements with pinhole sizes down to $50\ \mu m$. However, this size was difficult to use even when all of the alignment issues had been resolved (the smaller the pinhole, the more critical the alignment into it). Before we had seen the first single ion signal, the exact alignment procedure was yet to be determined and we needed a bigger pinhole rather than a smaller one to alleviate the search problem.

Another way to reduce the amount of scattered light that reaches the detector is to improve the output optics. As discussed in Chapter 4, the current apparatus uses a mirror to collect ion fluorescence and focus it into the detector. This arrangement causes two types of aberration. First, since the PMT must be placed below the trap (so as not to block the fluorescence) there is astigmatism due to off axis propagation [Bor80, p. 211]. Second, due to the size of the vacuum chamber surrounding the trap, the distance from the trap to the mirror and that of the mirror to the detector, are not the same. This, taken together with the off axis propagation, leads to

coma [Bor80, p. 211]. Both types of aberrations cause unwanted scattered light to be focused into the pin hole. One solution, of course, is to use a corrected lens. This can be quite expensive, especially for a UV transmissive lens. However, such a lens will be used in the future on the Indium experiment.

Another approach to background reduction is to reduce the amount of light that is scattered. This can be done by changing the type of trap or its geometry, or by improving the shape and size of the focused laser spot. Increasing the trap size is not particularly practical since very high voltages would be required to maintain the same trapping parameters. Improving the characteristics of the focused spot is a good idea in general and we spent a great deal of time trying to do so on the Indium experiment. Even though we were able to get some improvements, they weren't dramatic (we needed at least order of magnitude improvements and we were getting at best 50%). Improving the beam shape at the focus remains an open issue for this experiment and one that should continue to be addressed. However, it occurred to me that we might be able to modulate the laser light and take advantage of the ion's finite lifetime to greatly reduce the background. This turned out to be true and was critical to our success in detecting the weak single ion signal. The technique that we used came to be called background suppression and is the subject of this chapter.

5.1 Overview of the Background Suppression Method

By taking advantage of the finite lifetime of the In^+ cooling transition, it is possible to eliminate the background by alternately turning the incident light beam on and off and then only turning the detection on when the light is off. When the light is off the signal can still be detected because of the transition lifetime but the background associated with the light is gone. The total amount of signal detected per unit time is reduced (because it is decaying rather than steady and the detection duty cycle is less than 100%), but since the background is not present, the signal to background ratio may go up dramatically along with an improvement in the signal to noise ratio. The amplitude modulation of the laser has a half period comparable to the resonance lifetime of the In^+ cooling transition and the detector is on only when the laser is off. The lifetime of this transition is 440 ns [Pei94] which makes it possible to use relatively simple electronics to produce the chopping signal and detection gating. The modulation is accomplished by placing an EOM in the second

harmonic generation enhancement cavity and applying a square wave voltage to it. This essentially changes the optical path length of the cavity (at the frequency of the square wave) thereby taking the cavity in and out of resonance with the fundamental. When it is on resonance, UV is produced. When it is off resonance, UV is greatly reduced. Meanwhile, the detection is gated so that it is only on when the cavity is off resonance, producing very little UV. With this modulation in place we have observed a reduction in background by as much as a factor of 500 and an associated increase in the signal to noise ratio by a factor of 5.

5.2 *Theoretical Analysis of a Two State System Response to Square Wave Modulated Excitation*

We consider here a two level atomic system with excited state lifetime τ . The atomic system is excited by a laser which is chopped on and off with square wave modulation. The duty cycle of the modulation is not constrained to be 50%. With the laser power near saturation the excited state population is nearly 50% at the end of one laser-on pulse. During the subsequent detection time, τ_d (and corresponding laser-off time), most of the atomic system decays so an appreciable fluorescence can be detected. Since the laser is off during this time the background is almost reduced to nothing (see Figure 5.1).

To determine the optimal modulation period and duty cycle (the fraction of the modulation period during which the detection is on) it is useful to calculate the response of the system to a square wave modulation of the laser. To a first approximation this can be achieved by analyzing an RC circuit driven by a square wave modulated voltage. The results are essentially the same as those obtained by solving the two state rate equations and are particularly useful in the case where the laser is broad. However, if the laser is narrow (where narrow means that its coherence time is long compared to the Rabi period of the two level system) then coherence effects start to become important and a more general approach using the two level optical Bloch equations in the Rotating Wave Approximation [Lou83, p. 54] is appropriate. Taking into account the finite laser width, these equations are:

LASER ON:




Figure 5.1: Schematic showing UV light modulation and detection gating.

$$\frac{\partial \rho_{22}}{\partial t} = -\frac{i}{2}V^* \exp(i(\omega_0 - \omega)t + \delta\phi) + \frac{i}{2}V \exp(-i(\omega_0 - \omega)t + \delta\phi) - 2\gamma\rho_{22} \quad (5.1)$$

$$\frac{\partial \rho_{12}}{\partial t} = \frac{i}{2}V \exp(-i(\omega_0 - \omega)t + \delta\phi)(1 - 2\rho_{22}) - \gamma\rho_{12} \quad (5.2)$$

LASER OFF:

$$\frac{\partial \rho_{22}}{\partial t} = -2\gamma\rho_{22} \quad (5.3)$$


$$\frac{\partial \rho_{12}}{\partial t} = -\gamma\rho_{12} \quad (5.4)$$

where ρ is the two level density matrix, V is proportional to the dipole matrix element


between the two states, ω_0 is the transition frequency between the two states, ω is the laser frequency, $\delta\phi$ is a random phase shift between 0 and 2π [Cut66], and γ is equal to $1/2\tau$. The quantity $\delta\phi$ has a uniform distribution and is applied at a random Poisson distributed time with mean equal to the laser coherence time to simulate the finite laser width.

These equations can be solved exactly. While the solutions to equations 3 and 4 are just simple exponentials, the solutions to equations 1 and 2 are quite complex especially when one accounts for arbitrary initial conditions. The complexity of the solutions make them difficult to use, so we numerically integrate the equations with a 5th order adaptive step size Runge-Kutta algorithm [Pre92].

This analysis was carried out in the specific case of the $5s^2 \ ^1S_0 - 5s5p \ ^3P_1$ transition of In^+ . The upper graph in Figure 5.2 shows the calculated excited state population as a function of time for a given modulation period and duty cycle assuming an ideal laser (infinite coherence time) with $100\mu W$ of power focused to a $40\mu m$ spot giving an intensity of $6.25W/cm^2$. The ion is assumed to be in an orbit smaller than the focused laser spot, hence it is always being irradiated by the laser. This is a good assumption for single laser cooled ions. The oscillations during the laser-on time are due to Rabi flopping and suggest that if the laser-on time, τ_l , is a half integral multiple of the Rabi period then maxima in fluorescence will be observed. Similarly if τ_l is an integral multiple of the Rabi period then minima in the fluorescence will be



(a)



(b)

Figure 5.2: Simulated excited state population as a function of time: (a) several laser on/off cycles and (b) expansion of the first cycle.

observed. Figure 5.3 shows the calculated fluorescence as a function of modulation period for several duty cycles. The expected maxima and minima are apparent along with an increase in fluorescence as the duty cycle is increased. The latter is explained by noting that the population inversion due to Rabi flopping is greatest during the first “flop” (spontaneous emission destroys coherence making subsequent inversions smaller). A larger duty cycle means a shorter τ_l corresponding to earlier inversions. Thus, at the beginning of the laser-off time the excited state population is higher, leading to more fluorescence. Figure 5.4 shows the same calculated fluorescence as a function of modulation period, but now with a 10 MHz wide laser modeled by random phase shifts as discussed previously. The same behavior of higher fluorescence for shorter τ_l is evident in spite of the added laser noise. This is not surprising since this coherence effect takes place on the time scale of less than 50 ns whereas the laser coherence time is 100 ns. Note that in Figure 5.4 (c) I have indicated points of local maximum and minimum fluorescence. The reason for these extrema has been described, but Figure 5.5 helps to demonstrate it visually (note that we are once again assuming an ideal laser for simplicity since this doesn't change the result). Graph (a) of this figure corresponds to a $.38 \mu\text{s}$ modulation period ($5\times$ the Rabi period) whereas graph (b) corresponds to a $.45 \mu\text{s}$ modulation period ($6\times$ the Rabi period). In (a) we see that since the laser is being shut off at the maximum population inversion (τ_l equals a half integral multiple of a Rabi period since in this case there is a 50% duty cycle), the excited state population (and hence the fluorescence) stays higher than in (b) where the laser is turned off at the minimum population inversion (τ_l equals an integral multiple of a Rabi period).

From Figure 5.2 (b) it is clear that the optimal τ_l is $1/2$ Rabi period. Figure 5.6 shows calculated fluorescence vs. τ_d with τ_l set at $1/2$ Rabi period. This graph predicts an optimal τ_d of 350 ns giving a fluorescence over 90% of what it would be if there were no modulation. This corresponds to a modulation period of approximately 388 ns and a duty cycle of 90.3%. An ideal laser with infinitely long coherence time was assumed.

Solving the equations with a coherence time of 100 ns simulating a $2\pi\times 10$ MHz laser width does not appreciably change the results since the flopping is happening on a shorter time scale. In addition, note that while the inversion level dies out during each laser-on time, it comes back up to its initial value after the laser has been off.




Figure 5.3: Simulated data for modulated fluorescence as a fraction of unmodulated fluorescence plotted against modulation period for several duty cycles: (a) 30%, (b) 40%, (c) 50%, and (d) 60%.






Figure 5.4: Calculated fluorescence as a function of modulation period for several duty cycles and a finite laser width of 10 MHz modeled with random phase fluctuations. The duty cycles are: (a) 30%, (b) 40%, (c) 50%, and (d) 60%.



(a)



(b)

Figure 5.5: Simulated data for excited state population vs. time for the two modulation periods, $0.38 \mu\text{s}$ and $0.45 \mu\text{s}$. This comparison helps to explain why one modulation period gives more fluorescence than the other (see text).




Figure 5.6: Simulated data for modulated fluorescence as a fraction of unmodulated fluorescence plotted against laser-off time (which is the same as the detection time, τ_d). Laser-on time, τ_l , is fixed at 37.5 ns (1/2 Rabi period). The value of τ_d for which fluorescence is maximized is indicated. The small periodic ripple is due to aliasing.

This is due to the fact that during the laser-off time, the excited state population is allowed to drop well below what would normally be the steady state level. Since the two level system behaves almost exactly like an underdamped harmonic oscillator, as long as the excited state population is below the steady state at the beginning of the laser-on time, it will always overshoot beyond the steady state value when the laser is first turned back on. Hence, the larger fluorescence obtained by taking advantage of the Rabi flopping is not a transient effect and is present for macroscopic times as long as the laser is being modulated.

It is interesting to compare this result with a similar calculation using the rate equations. A rate equation calculation predicts an optimal fluorescence which is only 32% of the value obtained by leaving the laser on continuously. This is much less than the 90% optimum predicted by the Bloch equations approach which should come as no surprise since the rate equations will never show coherence effects such as the Rabi oscillations that we have just taken advantage of.

In practice, our current electronics limits us to 400 ns pulses so we are not able to take full advantage of the Rabi flopping. Figure 5.7 shows fluorescence vs. τ_d with τ_l held fixed at 400 ns. It is clear that a τ_d of about 568 ns gives maximal fluorescence of 43.7% of the unmodulated fluorescence. A similar calculation including the finite laser width (finite coherence time) gives a maximal fluorescence of about 35% occurring at $\tau_d = 440$ ns as shown in Figure 5.8. Note that this is in agreement with the rate equation result because our pulse is now substantially longer than the coherence time. This combination of parameters has been implemented and the predicted fluorescence observed.

5.3 Background Suppression Apparatus

As was mentioned in the beginning of this chapter, the amplitude modulation of the laser is accomplished with an intracavity EOM (Gsänger model PM 25) in the SHG enhancement cavity. The high bandwidth of this EOM makes it possible to generate the RF modulation required to make this method work for allowed atomic transitions. A 70 volt amplitude of the modulation suffices to shift the cavity off resonance by 4 to 5 half widths of the cavity fringe which is roughly a Lorentzian line shape. Since the shift is occurring on the time scale of microseconds and the locking electronics has only a kilohertz bandwidth, the shift happens too fast to cause the cavity to go




Figure 5.7: Calculated fluorescence vs. detection time (τ_d) with laser-on time (τ_l) fixed at 400 ns. The value of τ_d for which fluorescence is maximized is indicated. An ideal laser with infinite coherence time has been assumed. The small periodic ripple is due to aliasing.




Figure 5.8: Calculated excited state population vs. laser-off time, τ_d , with a laser width of 10 MHz. The laser-on time, τ_l , is held fixed at 400 ns. The graph shows that the maximum excited state population on each cycle is reached for $\tau_d = 440$ ns.


out of lock. This shift is enough to reduce the circulating visible power by a factor of 22 and the generated UV power by a factor 500, which is also the background reduction factor. The modulation voltage is an amplified TTL signal whose duty cycle, frequency and phase relative to detection are all adjustable.

To set the modulation parameters we first set the laser-on time, τ_l . We then purposely misalign the input laser beam onto one of the trap electrodes to create an abnormally large scattering signal. Next, we adjust the relative phase of detection to minimize this background. To get the correct duty cycle, we note that the number of detected counts will come down sharply as the duty cycle is reduced, if the detection time is overlapping the laser-on time. Once the detection gate falls completely outside of the laser-on time, the detected counts will fall off much less sharply. Consequently, tuning the duty cycle to correspond to this change in the rate that counts are falling off as a function of duty cycle gives an optimal signal. In practice, several iterations of this procedure are required for best results. Note that this *fine tuning* of the duty cycle is not to be confused with the calculation of its optimal value in Section 5.2. The calculated value is used as a starting point. Fine tuning then takes into account the fact that the UV light does not shut off instantaneously, among other details.

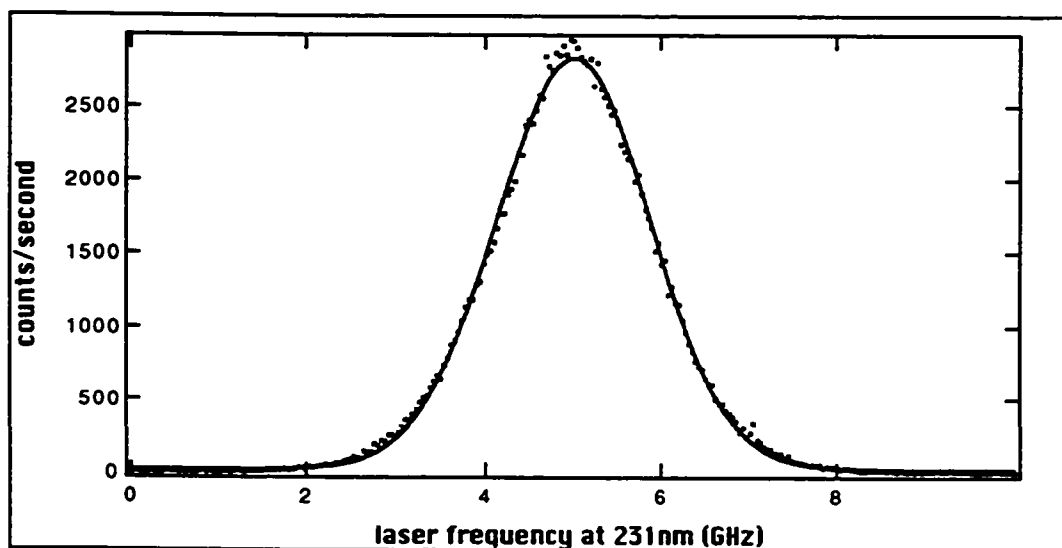
5.4 Background Suppression Results

With a 70 Volt modulation at a frequency of 1.1 MHz and approximately 50% duty cycle we were able to achieve a factor of 500 decrease in background and a factor of 5 increase in signal to noise (see Figure 5.9). While this is a nice improvement, it is not particularly useful for signals which are large compared to the background. However, this scheme becomes very useful when the signal is instead quite weak, and otherwise might be lost in the noise. This is demonstrated in Figure 5.10 which shows a weak In^+ ion cloud signal with and without the background suppressed. When background suppression is turned off, the signal is barely discernible in the noise.

As a preview to the next chapter, Figure 5.11 shows our first single ion signal with and without background suppression. In each graph the laser was blocked just after passing the center of the resonance to prevent laser heating (and possible expulsion of the ion from the trap), hence the sudden drop in fluorescence. Clearly, it would have been very difficult to detect this signal without some enhancement in the signal to noise ratio.

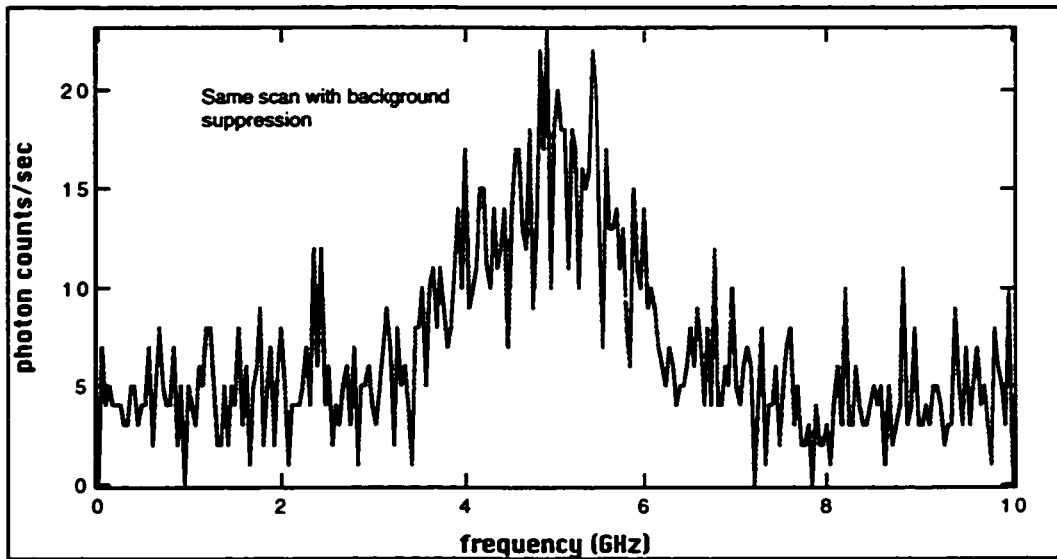


(a)




(b)

Figure 5.9: Strong fluorescence signal from a large ion cloud (a) without background suppression and (b) with background suppression. The scale in both graphs has been intentionally made the same so that the change in noise can be easily seen. Note that background is almost reduced to zero and the associated noise is also dramatically reduced.

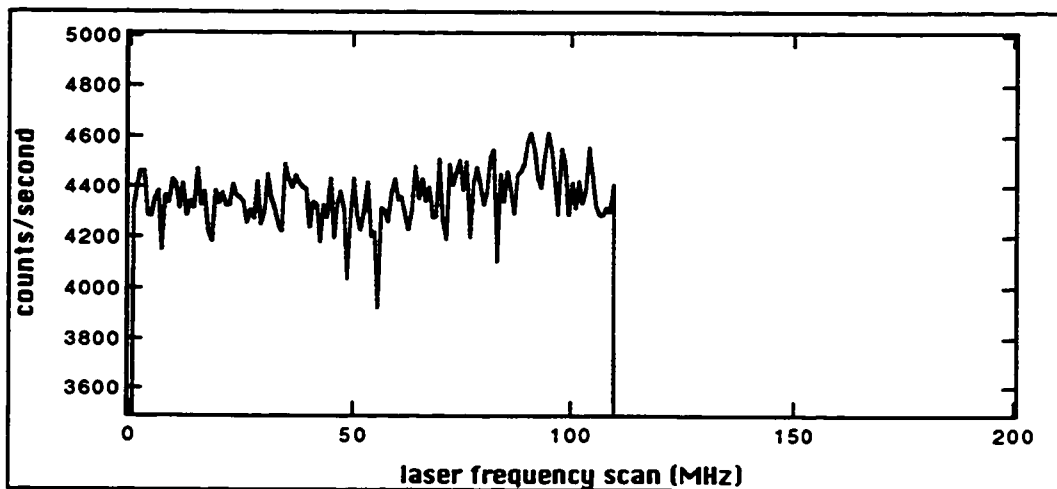


(a)

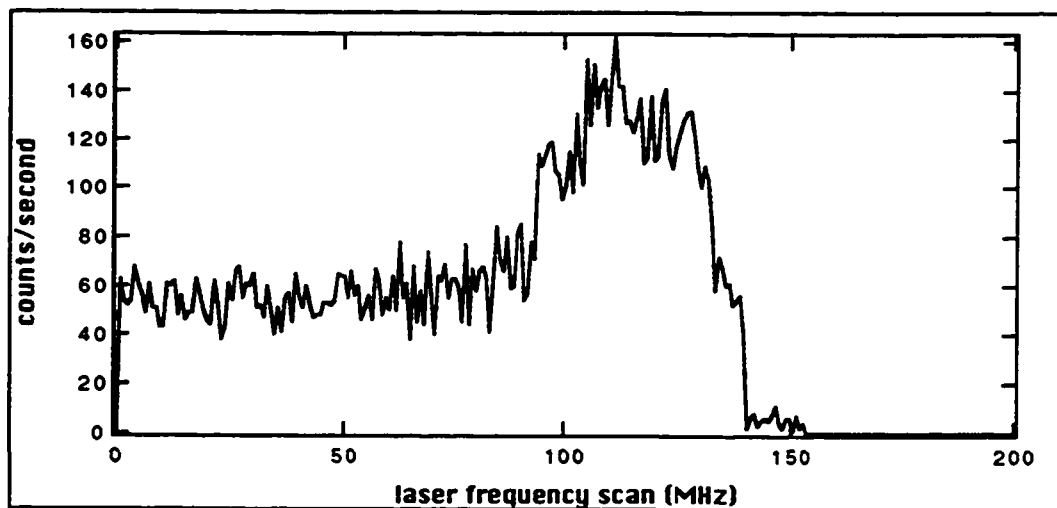


(b)

Figure 5.10: Comparison of a weak ion cloud signal (a) with and (b) without background suppression. The signal was weakened by reducing the number of ions in the trap.



(a)



(b)

Figure 5.11: Weak single ion signal from the In^+ cooling transition (a) with no background suppression (note that the sudden drop off in signal is due to the laser being blocked after the frequency has passed the transition so as to avoid ion heating) and (b) the identical signal with background suppression.

Chapter 6

RESULTS

In this chapter I will discuss the results of our single ion experiment. Our work started with electrical detection of large buffer gas cooled ion clouds, followed by optical excitation and detection of clouds and finally, single ions. The work with clouds proceeded quickly and because of the large amount of experience trapping other ion species in this lab we were optimistic that we would also be able to see a single In^+ ion signal quickly. Our initial optimism turned to concern however, as the search for a single ion signal went from weeks to months and finally to years. Single laser cooled In^+ ions proved difficult to trap due to the weak cooling transition. We knew this would be the case to some extent, but it would take some time for us to determine exactly how to allow for it. Initially, we used procedures that had been used with success on other species. When these didn't work we began to modify them. I will go into detail on what we tried and what finally worked.

6.1 Buffer Gas Cooled Ion Clouds

There are several factors involved in successfully trapping ions. First, a valid loading procedure must be determined so that ions can be reliably placed in the trap. Second, *a priori* procedures for optical alignment must be established so that both the laser beam into the trap and the fluorescence into the detector can be set up without the aid of an ion signal for optimization. Third, the correct frequency of the atomic transition to be studied must be known quite accurately (within a small factor of the transition width) so that the search in frequency space is minimized.

Large buffer gas cooled ion clouds were chosen for initial work for several reasons. First, buffer gas cooling (thermal dissipation through collisions with the buffer gas—see Section 2.3.1) is very efficient and requires no alignment. Of course, the ultimate cooling limit is essentially the ambient temperature which may be many orders of magnitude greater than that achievable with laser cooling. Buffer gas cooling can only be considered an intermediate step. But it is a good one, because ultimate

cooling temperatures are not desired when one is searching in frequency space for a signal. The Doppler broadening of room temperature simplifies the search problem. In addition, the large number of ions makes the signal stronger and also easier to detect when other parts of the apparatus may not be optimized.

At the beginning, electrical detection was used because it is a well established technique and required neither the laser, nor its associated second harmonic generation, to be operational. This way we were able to focus on just the loading apparatus. Once we were able to reliably load ion clouds we could then turn to optical excitation and detection with confidence knowing that there was an ion cloud present to work on.

6.1.1 Loading Procedure

Initially the Helium leak was a sealed bulb with no capacity for changing the amount of Helium contained within it. It is desirable not to have too much Helium in the bulb because the leak rate (even when the heating coil is off) is proportional to the pressure in the bulb (see eq. 4.10). An attempt to load Helium into the vacuum chamber with only 2 mm of Helium in the bulb was unsuccessful. A pressure of 200 mm was found to work quite well and eventually even 100 mm worked fine.

The following is the loading procedure that we used in the initial stages of the experiment.

- **Turn off the pump for the vacuum chamber.**
- **Helium is loaded into the vacuum chamber** by applying a current through the heating coil with a Variac set to between 10-12 V. Later we would step down the Variac to get finer control on the coil voltage. The particular voltage setting was chosen to give the coil a dull red glow in room light (I didn't record the actual current). The heating coil is left on for about 5 minutes resulting in a Helium pressure in the vacuum chamber of about 10^{-5} Torr.
- **Turn on the Indium oven.** This was controlled with a Variac that was set to 42 V (roughly 4 A). The oven was allowed to warm up for about 1 minute.
- **Turn on the filament.** This is done concurrent with the oven warmup. It was initially controlled by a Variac and was turned up very slowly (about 30 seconds

or more to reach operating voltage) to avoid burning out the filament. Later, a constant current supply was used to operate the filament so that we could simply turn it on without having to ramp up the voltage. The filament voltage was set to about 30 V¹ to give 50 μA of electron beam current. Later, 25 μA of electron beam current was found to be sufficient. In Chapter 4 I mentioned that the electron beam current is implemented by changing the potential of the filament relative to the nearby anode. A potential of -190 V turned the electron beam on and $+50\text{ V}$ turned it off. A filament warmup time of about 30 seconds was required before the desired electron beam current would be generated.

- **Pulse the electron beam current on.** Once the filament and oven have warmed up, the electron beam current is pulsed on for several seconds (later we found that the ion signal would continue to increase if the electron beam was left on for up to 1 minute).
- **Turn off filament and oven.**

6.1.2 *Electrical Detection*

The ramp voltage applied to the ring electrode of the trap is also applied to the X input of an oscilloscope. After ions have been loaded, and with the electrical detection signal applied to the Y input of the scope, a slightly asymmetric resonance lineshape was displayed. The signal to noise ratio was about 50. The lineshape is roughly Lorentzian and its width is due primarily to the Q of the tank circuit being used to detect the signal. The asymmetry is due to the fact that the ion cloud is large and samples the anharmonicity of the trap (which will increase the further the ions are from the center). This anharmonicity could be exaggerated by changing the RF drive amplitude (Ω). A lower amplitude gives a shallower trapping pseudopotential. The shallower trap allows the ions to move in larger orbits where the potential is more anharmonic. As the RF amplitude was further decreased, the anharmonicity would increase, until the signal would suddenly disappear. Once the trap depth gets smaller than a certain value, the ions are simply expelled from the trap.

¹ Unfortunately, in the early stages of the experiment, I didn't record the filament current, only the voltage.

We also observed a double peak in the resonance. This is another measure of the anharmonicity of the pseudopotential. That is, as the DC voltage was scanned up, the observed peak was at a different point than when the DC voltage was scanned back down. This effect is explained very nicely in [Lan76, p. 88-89]. The basic idea is that as the amount of anharmonicity in the potential is increased (or equivalently as an ion's orbit increases causing it to sample more anharmonicity) the frequency at which a peak in the resonance lineshape occurs is "pulled" up. Eventually the pulling becomes so severe that a region of instability is created leading to a hysteresis in the lineshape. There are two stable solutions between which the ion jumps, and it jumps at different points depending on whether the frequency is being increased or decreased.

Using a very simple argument based on the energy density due to Coulomb repulsion of ions packed together, we estimate a cloud size of between 10^5 and 10^6 ions. Visual observations of ion clouds trapped in the Ba^+ experiment suggest that these clouds reach a size of approximately 1/3 of the ring inner diameter before leakage becomes significant. This corresponds to 1/10 the well depth. This estimate should be taken is *very* approximate, since it may not apply very well to the In^+ configuration (we can't see these clouds since they fluoresce in the UV). See Section 2.3.3 for more details.

Once loaded, the ion signal stayed for months at a time. Occasionally, it was necessary to pump out a small amount of the buffer gas. Due to leakage (quartz remains slightly permeable to Helium at room temperature), the pressure would tend to go up slightly on the time scale of days or weeks. The leakage was apparent in the increased width of the resonance line. As the Helium was pumped out, the line would become sharper at first until too much Helium was pumped and buffer gas cooling was no longer strong enough to balance heating effects. At this point, the signal would abruptly disappear. After some experience, care was taken to turn off the pump before the ion cloud was ejected from the trap.

6.1.3 Optical Alignment Procedures

The optical alignment procedure for ion clouds was relatively straight forward. First, the beam is focused into the trap by using the combination of the two input lenses. The corrected lens is placed as close to the vacuum chamber window as possible and

the other lens is adjusted to get the resultant focus at the trap. The focused beam is moved within the trap by a 2-D translation stage holding the corrected lens. The beam is observed on the opposite side of the vacuum chamber with a fluorescent card. It is focused when the amount of translation required to eclipse the beam with the trap ring electrode is minimized. When the beam is focused right at the trap, its eclipse appears to come from all sides symmetrically (instead of from just one side, as is the case when the focus is not at the trap).

Next, the output focus must be set. To do this, the input beam is translated so that it is hitting the trap electrode. This creates a UV point source relatively close to the trap center and focusing this into the PMT will be very close to focusing an ion light source into the PMT. Fluorescent paint is placed on the pin hole covering the PMT so that the point source can be observed there. The position of the PMT is moved slightly so that the point source image is just to one side of the pin hole. The PMT is then moved along the axis of the output mirror to find the best focus of the spot on the fluorescent paint.

To remove some of the astigmatism in the output fluorescence, a correction plate is tilted in front of the PMT. As has already been described, this corrects some of the astigmatism at the expense of shifting the position of the fluorescence on the PMT. When the trap is back-lit, its shadow appears on the pinhole. We calculated the difference between the shift of this shadow and the shift of the UV to be about 350 microns (see appendix B.1). Once an ion signal was obtained we were able to verify that this calculated difference was correct. Consequently, the output alignment procedure was to center the pinhole in the shadow of the trap, and then shift it to account for the correction plate refraction. This was in the vertical direction since the plate was tilted around a horizontal axis. There was also a smaller deviation from center in the horizontal direction due to irregularities in the trap geometry and to how these are imaged in the trap shadow.

Similarly, the input alignment is accomplished by centering the focused spot in the trap. The spot is translated from one side of the trap to the other (visible only as it eclipses the focused beam). The indicated distance between the two sides of the trap is determined and divided by 2 to get the center location. This is done in one direction, then in the orthogonal direction and then repeated for better accuracy. As with the output alignment, once we had an ion signal we were able to develop

empirical adjustments to this centering algorithm.

6.1.4 Optical Fluorescence Detection

While making our original search for an optical ion cloud signal, we were aided by the electrical detection system that told us our loading procedure was working. Since the cloud was expected to take up a substantial portion of the trap volume, input alignment was not critical. Likewise, the input and output foci, though important, were probably not critical. We expected to be able to use *a priori* alignment procedures and be close enough to see a signal. We had hoped to make output alignment also non-critical by using a large PMT pin hole, but even without the high power (and associated background) delivered by a doubling enhancement cavity (our initial work was done with only single pass doubling through our BBO crystal), the background still proved to be a limiting factor. Consequently our search boiled down to two areas: output alignment and laser frequency.

At the time of our initial search the only source of information we had on the correct wavelength of the $In^+ \ ^1S_0 - ^3P_1$ cooling transition was the comprehensive study of Paschen and Campbell in 1938 [Pas38]. This study gives a wave number of 43350.62 cm^{-1} for the $F = 9/2 - F = 9/2$ hyperfine (HF) line. We split off part of our visible beam into a Burleigh Wavemeter to monitor our exact laser frequency and searched within plus and minus 1-2 cm^{-1} from the Paschen value.

To alleviate the output alignment search problem we started with a 200 μm pinhole. This gave us about 200 s^{-1} background with about 1 μW of UV power incident on the trap. After several days of fruitless searching in frequency space, we decided to try a 100 μm pinhole to reduce the background. We now had only 25 s^{-1} of background and saw the signal almost immediately. The initial signal was only 35 s^{-1} which explains why it was difficult to see it with the higher background. The wavemeter indicated a central wave number for the resonance of 21675.15 cm^{-1} or 43350.30 cm^{-1} in the UV. We suspected that this was the $9/2 - 9/2$ line but were concerned about the discrepancy between our wave number and Paschen's. To verify that this was indeed the $9/2 - 9/2$ line we searched for the other two hyperfine lines expected to be about .4 cm^{-1} above and below this one in the visible. Both were found immediately but the slight discrepancy from the Paschen values continued. Our value for the $9/2 - 7/2$ transition was 43349.24 cm^{-1} compared to the expected

value of 43349.56 cm^{-1} and for the $9/2 - 11/2$ transition our value was 43351.60 cm^{-1} compared to the expected value of 43351.90 cm^{-1} . In each case we were low by $0.30 - 0.32 \text{ cm}^{-1}$. Later we would find out that work done concurrently by Peik et al. [Pei94] agreed completely with our results. As expected, the largest number of counts came from the $9/2 - 11/2$ HF line (because of the larger number of decay pathways). On this transition we saw 63 s^{-1} with 20 s^{-1} background.

From eq. 3.2 we expected to see roughly 260 counts per second for each μW of incident power. Now that we had an ion signal, and could optimize our alignment, we were able to improve our signal towards this goal. Eventually, we achieved a signal strength of 500 s^{-1} with 245 s^{-1} of UV background (the UV power was $5-10 \mu W$ out of the crystal or about $2-3 \mu W$ at the trap) on the $9/2 - 7/2$ HF line. This signal is shown in Figure 6.1 along with a fit to a Gaussian lineshape. A power of $2 \mu W$ at the trap should produce about 520 counts per second which is in quite good agreement with the observed value. Using the notation of eq. B.5 we obtained in this case: background = $k_0 = 250 \text{ s}^{-1}$, signal = $k_1 = 497 \text{ s}^{-1}$ and $1/e$ half width = $k_3 = 24$ channels. Using eq. B.6 this value of k_3 corresponds to a full width at half maximum in the UV of $\delta_{uv} = 1.56 \text{ GHz}$ and using eq. B.7 this gives a temperature of $T = 325^\circ \text{K}$ — as expected, just above room temperature. This signal was obtained during the early stages of the experiment before the background suppression system had been implemented. Later, with the second harmonic generation enhancement cavity producing several hundred μW of UV power and the background suppression scheme in place we were able to get a signal like the one shown in Figure 6.2 for the In^+ $9/2 - 11/2$ HF line. In this graph the signal to noise ratio is good enough to show up the discrepancy between the Gaussian fit and the actual lineshape (Voigt profile). The data does not rise or fall quite as rapidly—characteristic of a Lorentzian influence.

6.1.5 Tellurium Saturation Spectroscopy

The wavemeter used to monitor our laser frequency was very convenient, but ultimately we wanted to lock the laser frequency against drift in addition to simply knowing what its frequency was. To this end, we incorporated a Tellurium saturation cell. As discussed in Section 4.5, the idea was to lock the laser to a molecular line to prevent drift. As a bonus, the pattern of molecular lines in the vicinity of the locking




Figure 6.1: Resonance signal from the $In^+ 1S_0 - 3P_1 F = 9/2 - F = 9/2$ transition. This signal was gathered with approximately $2 \mu W$ UV power incident on the ion cloud, a $100 \mu m$ PMT pinhole, and 1 s of integration per channel. Each channel corresponds to 40 MHz in the UV. This spectrum was taken without background suppression or the SHG enhancement cavity.




Figure 6.2: Ion cloud signal with both SHG enhancement and background suppression.

line provided a “road map” enabling us to locate the correct laser frequency at the beginning of the day *without* the wavemeter.

In order for locking to be feasible, we needed to have a Tellurium molecular line within 200 MHz and preferably 100 MHz of the In^+ line. This facilitates the use of inexpensive AO Modulators to shift the laser frequency from the In^+ line to the Tellurium line. Our initial search for an appropriate molecular line was performed with a single-isotope Tellurium cell. With this cell we were unable to find a line within 500 MHz of any of the In^+ HF lines. Next, we tried a naturally occurring mix of Tellurium isotopes and were rewarded with a molecular line within 100 MHz of *all three* of the In^+ HF lines.

The next step was to determine, as accurately as possible, where each of the potential molecular locking lines were relative to their corresponding In^+ HF lines. For ion clouds this was not a concern since the width of the ion cloud resonance (~ 1.5 GHz) was already much larger than the frequency difference between the In^+ line and the molecular line (< 100 MHz). However, for laser cooled single ions where the resonance width is less than 400 kHz, precise knowledge of the shift was required. A typical curve fit of 6 GHz scan over the In^+ resonance gave a one sigma error in the center location of about 7 MHz.

The first technique used to determine the shift was to scan the In^+ line, switch the photon counting input going into the computer from the In^+ fluorescence to the Tellurium saturation signal, and then rescanning the same frequency range. The two scans could then be compared. In this approach, the two scans were taken one after the other, which means that added to the already existing uncertainty of 7 MHz was the uncertainty due to the drift of the laser during the scans. Given that two complete scans at 1 second/channel integration take about 12 minutes, and that the laser (commercially stabilized) drifts about 1 MHz/sec, this drift represented an additional uncertainty of 12 MHz. Adding the two errors in quadrature gives a frequency shift uncertainty of $\delta f \approx 14$ MHz.

To reduce part of this uncertainty, we reconfigured the interface board in the computer to multiplex the ion signal and the saturation signal together, so that each channel for each signal is being integrated at virtually the same time. This eliminates drift uncertainty. Another potential systematic was laser *power* drift. During the course of a scan, the laser power usually degraded by several percent. This introduces




Figure 6.3: Block diagram of electronics used to map the position of Tellurium saturation lines relative to In^+ HF lines.


a slight asymmetry in the resonance line shape and will introduce an error in the fit if not corrected. We chose to multiplex a third data stream into the computer consisting of the signal from a PMT measuring the strength of the UV beam used to excite the ions. The fluorescence signal was then normalized to the laser power using this third signal, eliminating power drift errors. The apparatus used is shown in Figure 6.3.

The UV beam signal used for normalization and the saturation signal are multiplexed together and sent to the computer on the "Photon Counts A" port. The ion fluorescence is input to the computer on the "Photon Counts B" port. The interface board has a digital countdown timer (on the 6822 VIA chip) which is used to control the multiplexer through the "MUX CTL" port. The same timer is used to demultiplex the signal on the board. The laser frequency is controlled from the "Laser Scan"


port. A software program was written that transfers the three data streams to three separate files. A dual scan of the In^+ $9/2 - 7/2$ HF line and associated Tellurium lines produced by this system is shown in Figure 6.4. Figure 6.4 (a) shows the normalization signal with a clear power drift. Figure 6.4 (b) shows the In^+ $9/2 - 7/2$ HF line fit to a Gaussian. Figure 6.4 (c) shows the Te saturation signal that was obtained using the first down-tuned sideband generated with a 92 MHz modulation. For this particular run, the center of the In^+ line was 0.9 ± 1.1 channels or 21.1 ± 25.7 MHz below the large saturation line in the center of the bottom graph.

Sidebands 1 through 4 down-tuned were analyzed in this fashion. For each sideband, 4 dual scans were taken. This made it possible to graph the frequency difference between the In^+ line center and the saturation line as a function of sideband number (all with the modulation frequency fixed). A straight line fit to this data gives a zero crossing, or point where the AOM induced shift makes the In^+ line and the saturation line coincide. The zero crossing will, in general, not correspond to an integral value of the sideband number. However, given the modulation frequency, we are able to extrapolate to the frequency shift and therefore the modulation frequency, required if one of the sidebands is to be used. The data is summarized in Table 6.1 and graphed in Figure 6.5. From Figure 6.5 we see that the zero crossing occurs at 68 MHz down-tuned. Therefore, for the In^+ $9/2 - 7/2$ HF line, the closest Te saturation line should be shifted into alignment by applying 68 MHz modulation to the AOM and using the second down-tuned sideband.


A similar dual scan of the In^+ $9/2 - 11/2$ HF line is shown in Figure 6.6. This one did not use the power normalization, but was done with background suppression for an improved signal to noise ratio. In this figure the 3 GHz scan refers to the visible laser frequency. The UV scan width is then twice this, or 6 GHz. A fit to the data indicated that the In^+ line was 27.1 ± 2.2 MHz below the specified Tellurium line with 100 MHz modulation being applied to the AOM and the second up-tuned sideband being used. This means that, given the center frequency determined from the upper graph, the frequency of this center was 127.1 MHz below the Tellurium line. This is due to the fact that the oppositely propagating beams in the saturation cell cause the saturation line to occur shifted by one half of the modulation frequency for each sideband. Thus, lines observed using the first sideband are shifted 50 MHz and those observed using the second sideband are shifted 100 MHz. Since we are



(a)



(b)



(c)

Figure 6.4: Dual scan of In^+ $9/2 - 7/2$ line and Te saturation signal. Graph (a) shows the normalization signal used to cancel errors due to laser power fluctuations. Graph (b) shows the In^+ line with the number of counts normalized. Graph (c) shows the saturation signal. The In^+ spectrum was taken without background suppression or SHG enhancement.

Table 6.1: Data for the frequency difference between the In^+ $9/2 - 7/2$ line and a Tellurium saturation line as a function of AOM sideband. The difference, δf , is measured in numbers of channels. In these spectra, each channel corresponds to 23.4 MHz. The negative sign in front of the sideband number indicates that this is a *down-tuned* sideband. The sign in front of δf indicates whether the In^+ line was above (+), or below (-) the saturation line.

<u>Sideband</u>	<u>Run</u>	<u>δf(channels)</u>	<u>Avg.</u>
-4	1	+12.8 \pm 1.3	
-4	2	+12.0 \pm 1.4	
-4	3	+9.7 \pm 1.3	
-4	4	+13.0 \pm 1.4	
			+11.9 \pm 0.8
-3	1	+5.2 \pm 1.5	
-3	2	+9.5 \pm 1.1	
-3	3	+9.1 \pm 1.0	
-3	4	+7.0 \pm 1.1	
			+7.7 \pm 0.7
-2	1	+1.1 \pm 1.2	
-2	2	+1.6 \pm 1.3	
-2	3	+2.4 \pm 1.1	
-2	4	+3.3 \pm 1.1	
			+1.7 \pm 0.7
-1	1	-0.9 \pm 1.1	
-1	2	-3.0 \pm 1.0	
-1	3	-1.9 \pm 1.1	
			-1.93 \pm 0.6




Figure 6.5: Frequency difference between the In^+ line center and the saturation line as a function of AOM sideband. The error bars are derived from four runs for each sideband.




Figure 6.6: Dual scan of In^+ 9/2 - 11/2 HF line without power normalization, but with background suppression. The In^+ spectrum is 3 GHz wide in the visible, or 6 GHz wide in the UV. The frequency where the In^+ transition reaches its peak is indicated in the Te spectrum. The AOM was modulated at 100 MHz and the second up-tuned sideband was used.

using the second up-tuned sideband, we are shifting the light up 100 MHz and then the figure shows that the Tellurium line was an additional 27.1 MHz higher or 127.1 MHz altogether. The error of 2.2 MHz is in the location of the Tellurium line relative to the In^+ line center. The error in determining the latter was 3.0 MHz. Therefore, the correct shift as indicated from this data was 127.1 ± 3.7 MHz.

When we finally saw our single ion signal, we were able to determine that the shift was really 94 ± 5 MHz. This was outside of the error bars by over 25 MHz. Looking back on this data now, it seems that one likely explanation for the discrepancy is the lack of normalization in fitting the $9/2 - 11/2$ data. Since each channel corresponds to 23 MHz in the UV, a one channel error in the fit could easily lead to the difference that we observed. Another possible explanation is that there was a small amount of laser cooling causing a slight asymmetry in the In^+ resonance lineshape. The asymmetry required to create the error would not be readily visible from our fits. The effect of laser cooling would be to make the line fall off more rapidly above resonance due to a small heating factor. This would tend to shift a symmetric Gaussian fit downwards predicting a lower center. Our calculated line center was indeed lower than that finally observed with single ions. Of course, laser cooling is a major effect in the line shape of single ion resonances. The difference is that with the single ions it was expected, whereas with the clouds we deemed it too small to be of concern. For future reference then, it is clearly important to normalize the data by the UV power and to be more careful with the fit. In spite of this possible error in judgment we were able to use the frequency information gained from this analysis to narrow down our search in frequency space for the single ion signal.

6.2 *Laser Cooled Single Indium Ions*

The value of the buffer gas cooled ion cloud signal was that it enabled us to fine tune much of the apparatus on a real signal that would be very similar to a single ion signal. This was necessary since we expected the search for the single ion signal to be very difficult. We needed to reduce the uncertainties in each of the parameters associated with loading and alignment to make the search space as small and manageable as possible. In this section I will describe the search and go into detail on some of the more important techniques that we tried. Finally, I will describe what ultimately worked.

6.2.1 *Single Ion Search Strategy*

One of the first concerns in our search was how large the expected signal would be and whether it would be detectable. Section 3.2.2 derives expressions for the expected number of counts for both a buffer gas cooled ion cloud and a laser cooled single ion. Eq. 3.3 gives a value of 661 s^{-1} per μW of incident power for the single ion case. With saturation power well under a micro-Watt we did not expect to see more than 500 s^{-1} in any case. Since this model had agreed quite well with the observed value in the ion cloud case there was some optimism that the signal would be observable in the single ion case as well. This model doesn't take into account background suppression, but in a worst case the reduction in signal from background suppression is about a factor of 5 or a signal of 100 s^{-1} . Since we were routinely getting less than 20 s^{-1} of background, a signal of 100 s^{-1} should have been readily observable.

Another concern was polarization rotation due to the Hanle effect. This is discussed in detail in Section 2.9. There I determined that for the ambient magnetic fields we have near our trap, a factor of 2 reduction in signal can result from not aligning the input polarization properly with respect to the position of the fluorescence detector. The basic rule of thumb of aligning the input polarization perpendicular to the line joining the ion and the detector (which is generally what we did) should result in no more than 20% signal reduction. In order to get the maximum fluorescence, it is necessary to rotate the input polarization by about 34° from the perpendicular just mentioned. We did not adjust the input polarization beyond the perpendicular since the marginal increase didn't seem worth the added complexity.

I have already mentioned that trapping ion clouds allowed us to reduce our search space to output alignment and laser frequency. Strictly speaking this isn't quite true. Our ion cloud work enabled us to determine that our loading apparatus worked, but we still didn't know how our procedures would need to be modified in order to reduce the loading process to just one ion. In addition, we knew from calculations, like those in Sections 2.4.1 and 2.4.2 that a hot single ion might require some form of initial precooling. How necessary would this be and how could we accomplish it? As far as loading goes, initially our approach was to simply lower the values for the loading parameters that worked with clouds. We brought the oven and filament up to the same current, but pulsed the electron beam on for only several seconds instead of the minute or so used with clouds. When this didn't work we systematically went through

all reasonable combinations of oven voltage, electron beam current, and electron beam on-time. Much care was also taken to set the input and output foci as accurately as possible. We were able to set the input alignment accurately, but as I've already mentioned, the output alignment was made difficult by the presence of the correction plate.

A typical day of searching involved carefully aligning and focusing the UV into the trap, setting the output alignment to the position of our best estimate and then scanning the loading parameters and laser frequency. My co-workers on the Ba⁺ single ion experiment² had good success in loading single ions with what they called "continuous loading". The idea in continuous loading is to save time by leaving the loading apparatus on for long periods of time, but at low levels. This was done with the laser detuned slightly below the expected transition frequency. Eventually an ion would load and the laser would immediately begin to cool it. Shortly after that a signal would appear. Sometimes they loaded more than one ion. This was usually evidenced by increased signal strength and increased micromotion sideband amplitude (I'll return to the subject of micromotion later on). A good example of multiple trapped ions is shown in ref. [Jan84, p. 74]. Several things could be done at this point to get a single ion. The laser could be tuned *over* the transition to cause a small amount of laser heating which might kick some of the excess ions out. The trap depth could be reduced allowing the hotter ions to escape. Even if the depth isn't reduced the hotter ions may still escape, leaving one ion remaining in the trap.

The alternative loading approach, called "pulsed loading", is more time consuming. This is particularly true in the initial stages of experiment before all of the parameters have been refined. Pulsed loading consists of turning on the oven and the filament as usual and then pulsing the electron beam for several seconds. After the electron beam is turned off, the laser is scanned to find an ion signal. Because of the expected weakness of our signal (several hundred counts per second with perfect alignment and probably less before an ion signal was available for optimization) we were not able to tolerate the high background associated with leaving the filament and electron beam on continuously. With Indium's weak cooling transition (approximately 50 times weaker than that of Barium) it was essential for us to reduce the background as much as possible. We had already put quite a bit of effort into back-

² H. Dehmelt, N. Yu, W. Nagourney, G. Janik and X. Zhao

ground suppression and the 1000 or so counts per second from the filament could easily swamp the Indium signal. The background from the filament is present regardless of whether the electron beam is on. To eliminate the background, not only did we have to turn off the electron beam, but we had to turn off the whole filament circuit. This added more time to the search procedure, since each time it was turned back on the filament required about one minute to heat up to the point where the electron beam current was stable.

6.2.2 *Single Ion Search Space*

In order to see a single ion, the following items needed to be correctly determined and so they were all part of the “search space”: Input/output alignment, input/output focus, oven voltage, electron-beam current (determined by filament current and filament circuit potential), electron-beam on-time, PMT pin hole size, laser frequency and initial laser detuning. Some of these had been refined already from the ion cloud work, but as searching went on yielding no results, all of them became suspect. The actual details of the search are not too enlightening. Suffice it to say here that a great deal of time was spent varying one parameter or possibly a pair of parameters while holding the others constant.

In addition to varying parameters, the experimental apparatus went through several evolutionary changes to improve power, reduce background and improve detection efficiency. For example, we started with simple Magnesium Fluoride coated Aluminum mirrors on our input optics which had a reflectance of between 60–80% (and which tended to degrade with time). Eventually, we went with dielectric coated mirrors for our wavelength with better than 99% reflectance. Several other optics were optimized as well by going to higher transmission UV quartz. The input focusing lens was upgraded to a UV transmitting corrected lens. We greatly improved our UV power by setting up the enhancement cavity for the SHG. We implemented the background suppression scheme to dramatically reduce the background. For awhile we did our searches on the $In^+ 9/2 - 7/2$ HF line because it was associated with the strongest Te line. However, it had two drawbacks. First, because of the fewer number of decay pathways, the $9/2 - 7/2$ line is the weakest of the 3. The strongest, $9/2 - 11/2$, produces about 30% more signal. Second, if the ambient magnetic field is oriented in just the right way (see Section 2.8), it is possible to have no signal at all

due to optical pumping. While the latter problem was unlikely, it was still a concern and in any case the former problem was motivation enough to switch our search to the other line which we eventually did.

6.2.3 *Chronology of Single Ion Search*

Our first attempts at finding a signal were actually performed on a trap and vacuum chamber that I haven't referred to yet, largely because it was not the apparatus which eventually produced a result. This first apparatus consisted of a stainless steel tube modeled after that used in the Ba^+ experiment. It had 6 ports; 3 with sapphire windows, one with a glass window, one for the trap header and one for the pump. The trap was of the Paul-Straubel variety, manufactured by drilling a hole in a flat wire strip with an electrical-discharge machine. This trap had already been used at one time in the Ba^+ experiment to trap single ions. Other than this, the apparatus was very similar to that already described. One problem with this type of vacuum chamber is the narrow aperture provided by the sapphire windows. These windows are required for UV transmission (the quartz windows available on Conflat flanges must not be baked higher than 200°C and we need to bake at 350°C) and yet require a special U-shaped mount to attach them to the surrounding flange. This mount limits the size of the window. Larger windows are prohibitively expensive and would have required a complete redesign of the vacuum chamber. Finally, the transmission of these windows was about 10% lower than quartz at our wavelength. Eventually we came to realize that a quartz tube with quartz windows fritted on (all built by our glass blower), that can be baked at better than 400°C , gave us more flexibility and better performance.

In spite of the disadvantages, this tube was and still is, quite successful on the Ba^+ experiment and so we tried it for Indium. After several months of searching, I began to notice that the windows were starting to get coated by Indium from the ovens. I was careful to direct the oven aperture away from the windows when I installed it, but the Indium doesn't come out in a straight beam, so eventually the windows will acquire a coating.

I took the opportunity of taking apart the steel tube, to put the quartz tube back on the table and further refine the loading and alignment parameters on an ion cloud. Another reason for going back to ion clouds was to try to reduce the gap between

loading clouds and loading single ions. By that I mean we tried to see whether we could get to a small laser cooled cloud [Neu78] starting from a buffer cooled cloud. The first approach was to get a buffer cooled cloud signal, pump out the buffer gas (the signal disappears) and then pump gas back in to see if the signal reappears. In fact it did, confirming our suspicion that at least for a short while, the signal disappears not because the ions are gone, but simply because they heat up and the signal is too broad to see.

Next, since we knew that ions stayed in the trap for some time even after the buffer gas was pumped out, we tried to use evaporative cooling as a way of bridging the gap between buffer gas cooling and laser cooling. The idea was to get a buffer gas cooled signal, pump out the gas and then reduce the well depth. The hottest ions would leave the trap, hopefully leaving only a small number of cold ions which could then be laser cooled. Although this seemed promising, I was never able to get a laser cooled signal out of it. When I loaded gas back in, I found that a small number of ions were indeed still there, but probably too many to be laser cooled. In fact if I had pushed harder on this it might have eventually worked. If there are initially 50,000 ions in the buffer gas cooled cloud and I reduce the signal by a factor of 100 there are still be 500 ions left, which are probably too many. I tried reducing the signal further and further, beyond where it could still be detected in hopes that I would reach a threshold where laser cooling would take over and the signal would return, but without success.

It turned out that when the buffer gas was pumped out, the signal didn't completely disappear as we first thought. If we took long integration times and did an on/off resonance comparison, it was still possible to see a signal. The signal would continue to get weaker over the pumping period, but at one point was still visible after several hours of pumping. We considered this to be our first vacuum In^+ signal, though it is unlikely laser cooling was having much of an effect. Figure 6.7 shows data taken after the gas had been pumped on for 5 hours. The places where the laser was placed on and off resonance are indicated. Our pressure gauge indicated 10^{-9} Torr when the signal finally became lost in the noise. This corresponds to a pure Helium pressure of about 10^{-8} Torr. Once the vacuum came down into the 10^{-9} region I tried sweeping the laser from below the transition. I wanted to see if I could get any laser cooling and increase the signal. Only once did I see any effect and it wasn't




Figure 6.7: A vacuum ion cloud signal obtained after 5 hours of pumping. The indicated pressure was about 10^{-9} Torr which suggests a Helium pressure of about 10^{-8} Torr. The reductions in signal near channels 10 and 30, are caused by the laser being manually tuned off resonance.

repeatable. While loading and observing very small clouds I noticed that sometimes I could load without even turning on the oven! We attributed this to Indium being deposited on the surface of the filament. When it and the electron beam were turned on we were emitting both Indium and electrons.

I did an analysis of different pin hole sizes to see which one gave the best signal characteristics. The pin hole sizes varied from $60 \mu m$ to $300 \mu m$ in diameter. Some were purchased from Melles Griot and others were homemade. To manufacture the homemade variety, we used an extra FEP and “drilled” the hole with the FEP mounted as the cutting tool on an electrical discharge machine. We found that we could make any size we wanted with a nicely uniform round edge. The largest signal is, of course, obtained with the largest pin hole. However, it is accompanied by larger background so that the signal to background (SBR) and signal to noise (SNR) ratios give more information. The hope is that, since the signal source is very nearly a point and the background source is spread out, as the pin hole is made smaller the background will be strongly reduced while the signal will only be slightly reduced. This is less the case with a cloud than a single ion since the latter is much more

localized. The best SBR was achieved with the 75 μm pin hole. This was true even though the signal was reduced by a factor of 5 or 6 by going from the 150 μm to the 75 μm as an example, and confirmed our hope that the signal was far more localized than the background. On the other hand, the SNR was best in the 150 μm pin hole. This is easy to understand. Assuming that the image of the fluorescing cloud on the pin hole is substantially larger than the pin hole, as the pin hole size is increased, the signal goes up quadratically. Assuming the noise is shot noise, then it goes as the square root of the background which is going quadratically with the pin hole size. So the signal is growing faster than the noise. This is almost definitely not true for single ions, where a larger pin hole will not lead to a larger signal. In fact, an ion cooled into the Lamb-Dicke regime will reside in a region much smaller than the wavelength of the light used to excite it (see Section 2.4.3 and [Win87]). In our case, this is on the order of nano-meters and well below the diffraction limit for a focused 231 nm beam. Theoretically the smaller the pin hole the better. However, I found that with pin holes smaller than 75 μm , searching for the precise alignment became exceedingly difficult. Beyond 150 μm the background began to go up faster with pin hole size (as we begin to overlap the source of it more) and no advantage in either SBR or SNR was obtained for pin holes larger than 150 μm . The conclusion was to use the 150 μm pin hole for the single ion search since it gave reasonable SBR and SNR performance and it eased the alignment search problem.

Since we were opening the quartz tube to put in the trap we decided to take the opportunity to modify our electron beam apparatus by adding field emission points (FEPs) to the existing filament. FEPs were desirable because they operate with high voltage and low current (rather than the other way around for filaments) and so produce no background. Jumping ahead a little bit, it turned out that the FEPs worked very well generating an electron beam current. We were able to successfully load buffer gas cooled ion clouds in the Paul trap with the FEPs. However, when we attempted to load single ions in a Paul-Straubel trap, the high voltage necessary to make the FEPs work (about 500 V) combined with their close proximity to the trap interfered too much with the trapping potential. Ultimately we used the traditional filament for loading of single ions.

We also decided to try coating the trap with Aerodag. Aerodag consists of very small flakes of Carbon suspended in an organic solvent. It is usually used as a

lubricant. When it is sprayed on a surface, the solvent completely evaporates (enough that it doesn't impede ultra-high vacuum) leaving only the Carbon. On the trap surface the Carbon can act as a sort of anti-reflection coating. Indeed, a reduction in reflection of incident light, and hence background, has been observed using Aerodag on the Ba^+ experiment.

When we started to use this trap we noticed immediately that our ion cloud signal was decaying away on a time scale of tens of seconds. The first thing we suspected was that somehow a charge was building up on some part of the trap header, such as the oven, creating a large contact potential. We were able to rule this out, but I attached a large "bleeder" resistor to the oven just to be safe. Next, we considered what had been changed in going to this setup. We weren't using the FEPs at all now so we doubted that they were causing any problem. That left the Aerodag.

We noticed that the decay rate was interrupted if the laser was blocked. In fact we were even able to observe this overnight. So the problem seemed to be related to the laser as well as possibly the Aerodag. We found that a photo-electron current of about 1 nA induced by the laser was being emitted by the trap. We verified that it was laser related by turning the laser on and off and observing the current. The energy of our UV photons is about 4.9 eV so they are in the right neighborhood to be able to generate photo-electrons from Carbon or stainless steel (this particular trap had stainless steel electrodes). However, the work functions of stainless steel and Carbon are comparable so why the sudden change? Perhaps there was a co-moving UV spot that hadn't been present before. This brought us back to the beam shaping issue. We tried to see if changing the shape of the input beam changed the photo-current. We were able to have an effect by changing the waist and by translating the position of the waist in the trap. In fact, we were able to observe the strongest photo-current by hitting the ring or the upper endcap with the beam. Hitting the lower endcap reduced the effect by a factor of 20.

To gain a better understanding of the photo-current characteristics of stainless steel and stainless steel coated with Carbon, I created a test system consisting of two endcap-shaped electrodes—one bare stainless steel and the other stainless steel coated with Aerodag. The test system was baked out and brought down to 10^{-8} Torr. Strangely enough, when the laser was directed onto the electrodes, both gave off the same amount of current.

Finally, we tried spatial filtering of the input beam. With a series of pinholes, and normalizing by the power in the beam *after* the pinhole, we found the ion signal decay rate greatly reduced with the smaller pinhole. Strangely, the photo-current was largely unaffected by the pinhole size.

As time passed, we noticed that the decay problem seemed to improve and, eventually, all but disappeared. One possibility was that there was some residual organic compound in the Aerodag that was heating the ions and causing the signal to disappear. After time this residue would eventually be pumped out by the ion pump. It was possible that with a heatable Paul-Straubel trap the situation would be improved, since we would be able to hasten the removal of the residue. This might partially explain why there were no decay problems observed on the Ba^+ experiment.

We spent quite some time loading small ion clouds and attempting to bridge the gap between buffer gas cooled clouds and laser cooled single ions. We decided to make the gap smaller by replacing the Paul trap (ring and endcaps) with the Paul-Straubel trap (ring only) that we wanted to use for single ions. The goal was to trap buffer gas cooled clouds, pump out the buffer gas and without changing any of the alignment, attempt to trap single ions. Thus we would use the cloud as our initial alignment, thereby eliminating much of the uncertainty associated with an *a priori* alignment procedure developed for one trap, but applied to another. In addition, having the ability to heat the trap might completely eliminate the decay problem *and* reduce potential micromotion problems.

The new trap that we used is the one described in Chapter 4: a twisted pair of Tantalum wire wrapped around a mandrill to form a ring (see Section 4.9). Making the trap out of thin wire instead of drilling a hole in a strip makes it possible to get appreciable heating of the trap by running a current through it. This is how we eliminate of micromotion-causing impurities that get deposited on the trap surface.

The RF drive voltage is connected to both ends of the Paul-Straubel trap so that no current runs through it during normal operation. However, to heat it we removed the RF drive and attached the trap electrodes to the oven electronics. In this way, the trap completes the oven circuit instead of the oven. The “oven” voltage is then brought up to 37 V resulting in a current of 2.5 A (enough to make the trap glow brightly). This is maintained for 5-10 minutes. The voltage is turned off and the trap and oven wiring is returned to its normal configuration. The first time we heated the

new trap we saw a dramatic change in pressure—from 1.2×10^{-10} to 3×10^{-7} . The pressure returned to normal shortly after the trap heating was complete. Subsequent heatings produced smaller pressure changes. Another side effect of heating was a slight repositioning of the trap due to temperature induced stress relief. Re-alignment of the laser was easily accomplished however and subsequent heatings had smaller effects.

With the new system still uncontaminated with Helium, I made an initial single ion search with only the rudimentary *a priori* alignment procedure determined from the other trap. After several days of no result I proceeded to load the Helium and look for a buffer gas cooled ion cloud signal in the Paul-Straubel trap (something that we hadn't done before). With the oven voltage set to 47 V (4.7 A—just enough current to make it glow dull red) and normal loading and alignment procedures determined from the previous trap, I found the buffer cooled In^+ $9/2 - 11/2$ HF line. The signal was weaker than in the other trap, due no doubt to the reduction in well depth in this trap. Decay of the signal was now almost undetectable. A fit to a full screen of fluorescence counts with the laser frequency fixed showed a decay rate of only 2-3 counts on a 60 counts per second signal over 5 minutes. This sort of decay rate is not easy to see with the noise level that we have and has probably been there all along. It is not a serious problem, but as an aside, it could have caused a slight asymmetry in the resonance line shapes that we used to determine the position of the In^+ line with respect to Tellurium lines. As the reader may recall these determinations were off by a factor well outside of the error bars (see Section 6.1.5).

After verifying that I was able to easily reproduce the ion cloud signal, I then set to work determining the optimal alignment parameters for this trap. Once I determined these parameters, I was able to demonstrate their accuracy. On subsequent days I would align the laser without an ion signal to use for optimization. Then I would load an ion signal and see how close the signal was to being optimal. I was usually within 15%. Now that I had more confidence in my alignment I went back to searching for the single ion signal again. After 3 more weeks, on July 7, 1994, I finally saw a laser cooled In^+ ion signal.

6.2.4 Loading and Observing Laser Cooled In^+ ions

When I first detected a laser cooled In^+ signal I was in the process of experimenting with the initial laser detuning. Also, given our negative experience with stray FEP voltage, on a suggestion from Prof. Nan Yu, I completely turned off the loading apparatus after loading. This eliminated the possibility of there being any stray DC biases around to interfere with the trapping potential. At first, this seemed unnecessary because such precautions were not required in the Ba^+ experiment. However, once again, the weak cooling transition of In^+ may require this sort of special handling and it certainly couldn't hurt. Later, once I had a signal, I was able to verify that the voltage due to the filament-off-bias (about 40 V) could indeed cause a problem, but wouldn't by itself always prevent the loading process from being successful. I was using the same oven voltage and electron beam current that I had successfully used with ion clouds. The only difference in loading was that I used a 2 second electron beam pulse instead of the much longer period for clouds. I set the laser scan width at 1 GHz (in the visible) and started it at 400 MHz below the estimated transition location. I was beginning to suspect that I was loading ions, but due to the weakness of the cooling transition, they were just too hot to see. So I tried keeping the laser frequency fixed at its initial detuning location for about 40 seconds before starting to scan up. This was not a new idea, but taken together with optimized power, background, optics, alignment, and the elimination of stray biases it was enough. As I scanned the laser up, I saw a broad "bump" at approximately the expected location. See Figure 6.8.

Once I found this signal and had verified that it wasn't a laser power fluctuation or some other systematic, I tried a narrower scan. Figure 6.9 shows a 100 MHz scan over the same signal. The laser is blocked just as the fluorescence begins to drop, to avoid laser heating. Note the narrower features on the broad top of the line. It was first thought that these might be Zeeman splittings but they are spaced about 10 MHz apart. At about 1.4 MHz per Gauss splitting in our roughly 1 Gauss field, Zeeman lines should be closer (and probably not resolved by our commercially stabilized laser). The 10 MHz separation suggests micromotion sidebands at the RF drive frequency. I will come back to this later.

The background suppression scheme that we developed was critical to our success as is clear from Figure 5.11. The first graph shows the same first signal that we

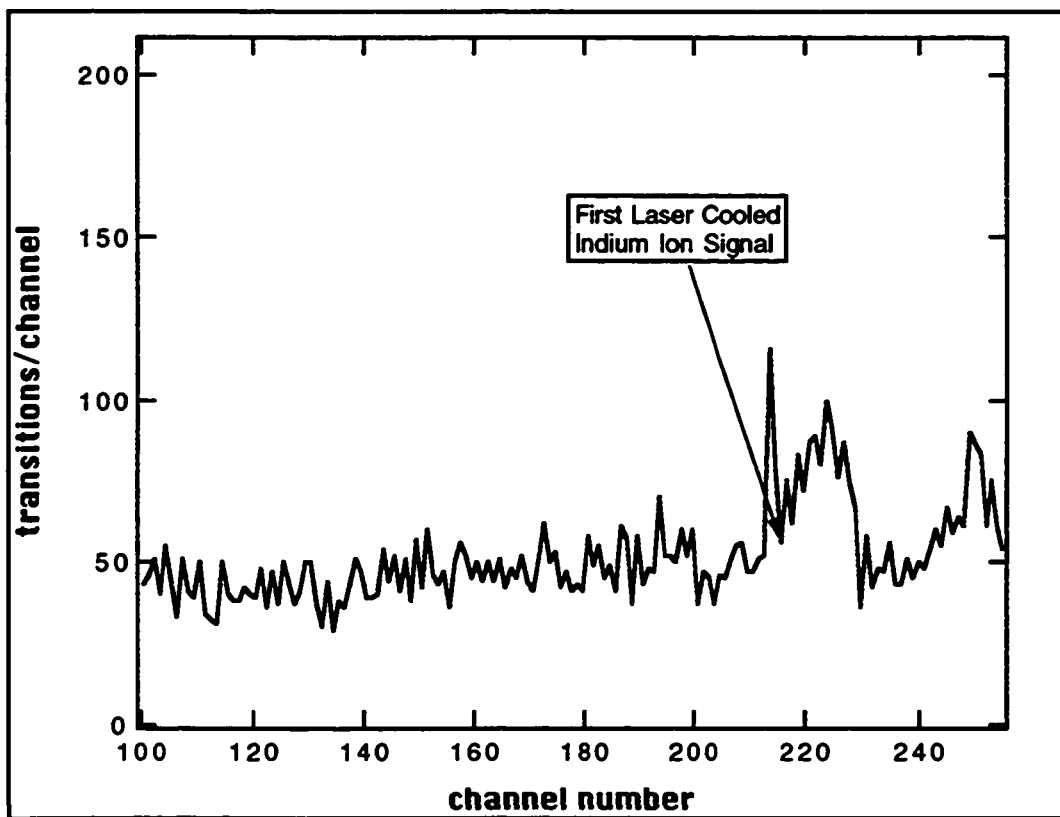


Figure 6.8: First laser cooled In^+ ion signal observed. The laser was being scanned manually at approximately 1 channel (about 8 MHz) per second. The second “bump” at the end of the scan is where the laser was backed off and then rescanned up.






Figure 6.9: First laser cooled ion signal with a narrower scan of 200 MHz. The laser is blocked just as the fluorescence drops (see text) accounting for the low count rate above the transition. The approximate HWHM is 20 MHz.

detected, but without background suppression. This weak signal would have been almost impossible to find without some sort of SNR enhancement. Now that we had the signal though, we were able to optimize our optics. This improved the signal enough that background suppression was no longer required to easily see it. Figure 6.10 shows another comparison of a signal without and with background suppression after the optics were optimized. Note also that the signal without background suppression, about 400 s^{-1} , is in very good agreement with the predicted value of 661 s^{-1} (see Section 3.2.2). This suggests that our understanding of the various loss factors and widths was good and that we can turn this around on the multi-ion case. There, our estimated signal agrees with the observed value *only* if we assume 50,000 ions in the trap. This is probably a more realistic number than the predicted capacity of 500,000 ions.

Optimizing the optics required locking the laser to the Tellurium line. This was necessary because the transition width was quite narrow (less than the laser width) which made it difficult to keep the laser frequency at the maximum fluorescence location. The 1 MHz drift per minute caused the fluorescence intensity to change substantially during the optimization process. Locking to the Tellurium line doesn't require special electronics since it is low frequency. A single pole low bandwidth integrator does the job. The output of our Tellurium saturation electronics box, which produces an error signal, is routed into the integrator and from there to the external control on the back of the 699-21 laser control box. Of course, when the laser is stabilized for high frequency excursions as well, the high finesse cavity is locked to the Tellurium line by feeding back the integrator output to the cavity instead of the laser. With the laser locked to the Tellurium line I was able to refine the *a priori* alignment procedure (see Section 6.2.9). The strongest signal that I was able to obtain in this way is shown in Figure 6.11. The narrowest line I obtained (without a stabilized laser) is shown in Figure 6.12. This spectrum was obtained after heating the trap to reduce micromotion (see the next section) and indeed, not only is the line narrower, but there doesn't appear to be any micromotion. A partial fit to the first half of this lineshape gives a FWHM of 10 MHz.



(a)



(b)

Figure 6.10: A comparison of a laser cooled ion signal after optical optimization, (a) without and then (b) with background suppression. In each case, a channel corresponds to 781 kHz and 1 second of integration. UV power was set to $50 \mu W$. A $75 \mu m$ PMT pinhole was used.




Figure 6.11: Best laser cooled ion signal. This one was obtained after optical optimization. Each channel is 781 kHz and 2.5 seconds of integration. The approximate HWHM is 9 MHz. The UV power was $50 \mu W$ and the PMT pinhole was $75 \mu m$.





Figure 6.12: Spectrum showing the narrowest laser cooled ion signal obtained without an externally stabilized laser. Each channel is 781 kHz and 1 second of integration. The approximate HWHM is 5 MHz.

6.2.5 Micromotion


We fit the optimized signal in Figure 6.11 to find its width. Since it is only the first half of the line (because I block the laser when the fluorescence begins to decrease), only the portion up to its maximum value was fit. This gives 12 MHz FWHM. In almost every scan that I did (including this one), I found that the fluorescence has two peaks: one right at the beginning and one just before I blocked the laser. The separation of these peaks was consistently between 5 and 10 MHz (since the laser width was comparable to this it was not possible to determine the separation very accurately). While the laser width prevents these peaks from being accurately resolved, the fact that they show up consistently suggests that they are part of the signal instead of noise spikes. Given the separation of roughly 10 MHz it seemed likely that the peaks were the carrier and lower sideband due to micromotion.

At the absolute center of the trapping potential, the field is zero. However, any distance away, the field is non-zero and oscillating at the RF drive frequency. The time average of this oscillatory field together with its inhomogeneity result in the time *independent* trapping pseudopotential. Away from the center though, the ion will see a time *dependent* RF field. Thus the RF drive causes a slow oscillation at the secular frequency (ω_i) and a fast (small) oscillation at the drive frequency (Ω). In our case the sidebands caused by the secular motion are at about 1 MHz and probably unresolved. Zeeman lines due to the splitting of the m_F sublevels in the ~ 1 Gauss ambient field are likewise of order 1 MHz apart (for the cooling transition). Consequently, the peaks that we observed are most likely due to the micromotion.

To test this, I waited until I had a signal that was broad and had well defined peaks. I then heated the trap and reloaded. Heating the trap is done to essentially clean it and reduce micromotion-causing contact potentials. The signal obtained after heating should be narrower and the peaks should be less well defined. Figure 6.13 shows an excellent example of the effect that heating the trap has on the signal. The overall signal is much narrower after the heating and the left peak is less pronounced, however the separation of the left peak and the central part of the line is different in each case. This would not be the case for micromotion sidebands. The resolution of this inconsistency will have to await a stabilized laser. Note that the consistent presence of micromotion sidebands is often an indicator of more than one ion in the trap. This issue will be taken up subsequently.



(a)



(b)

Figure 6.13: The effect of heating the trap on a laser ion cooled signal. Graph (a) shows the ion signal before heating and Graph (b) shows a different ion signal loaded just after heating. The improvement in micromotion is evident. In both graphs each channel is 781 kHz and 1 second of integration.

6.2.6 *Determination of the Correct In^+ Cooling Transition Frequency Relative to a Tellurium Line.*

Since the laser-limited linewidth of the ion fluorescence was now on the order of 10 MHz (at least after heating the trap) and the Doppler-free Tellurium lines were also about 10 MHz it was now possible to easily determine the relative frequency of the two within 10 MHz. We found that the correct detuning for the Tellurium line nearest the In^+ $9/2 - 11/2$ HF line was 94 MHz to the blue. I determined this simply by changing the modulation on the saturation cell AOM until I obtained maximum fluorescence. As mentioned in Section 6.1.5 this disagrees with our estimate derived from clouds. Most likely the discrepancy is due to fit errors caused by systematics such as signal decay or a non-Gaussian lineshape. The saturation AOM is modulated at 94 MHz and we use the second uptuned sideband.

6.2.7 *Single Ions?*

It was easy to make the assumption that we had a *single* laser cooled ion, but this needs to be proven. We only knew with certainty that we had a laser cooled cloud of no more than a few ions. The signal was not visible unless the laser was scanned up from the red which implies that the laser is cooling. As to whether it was a single ion or a small cloud of ions, more needs to be said. First of all, if it was actually a small cloud then experience with other ion species suggests that only a very small cloud of 2 or 3 ions is stable (Paul-Straubel trap). Any more than that and RF heating will cause them to leave the trap. Experience had also shown that small clouds did not persist in the trap for more than a few minutes without continuous laser cooling whereas a single ion could stay much longer. It was fairly common for the In^+ laser cooled signal to stay overnight. With no laser cooling for up to 20 hours, this is a strong indicator of only one ion. Our signal strength always seemed to have the same value of about 100 counts per second, also indicating one ion. If small clouds were being loaded, one would expect to see incremental amounts of signal for 1, 2 or 3 ions. Finally, I varied the loading parameters to determine the minimum required to load. Initially I loaded with the oven voltage set to 47 V. When I lowered it to 45 V, I would only load infrequently—perhaps once every 3 or 4 tries. This guarantees that loading is weak and it is unlikely that more than one ion at a time will be loaded. One other factor is that I *did* see signals with no resolved micromotion. This would be

highly unlikely with small clouds, however, this same apparent lack of micromotion could also be caused by a sudden laser frequency jump past the transition peak³.

Each of these arguments is compelling, but they don't constitute proof. If we are actually seeing 1 ion then it is suspicious that we don't *ever* see an indication of 2 or 3. One technique used to get just one ion from a small cloud, is to scan the laser just past the transition in order to cause a small amount of laser heating. This will kick out some of the ions. The signal should go down incrementally. It is disturbing that we have not been able to demonstrate this. Of course, the only way to prove how many ions we have in the trap is either by seeing some sort of incremental increase or decrease in signal, *or* to see quantum jumps.

If the ion jumps into a metastable state, the fluorescence will suddenly drop. This is what is referred to as a "quantum jump" [Coo85]. If it always drops to zero (the background level) then a single ion is indicated. Multiple ions will cause multiple levels of fluorescence to be observed as one or more of the ions is shelved in the metastable state. The demonstration of quantum jumps is the next target of this experiment.

6.2.8 Loading Parameters

With a laser cooled signal I was able to verify that certain extraneous voltages were lethal to the signal. Typically, my loading procedure involved turning off the entire loading apparatus (oven and filament) after the load was complete. If I then turned the filament voltage back on to its normal value the signal disappeared. What is more, the signal didn't return *at all* when the voltage was turned off, indicating that the ion(s) is(are) being kicked out, rather than simply displaced from the beam. Not surprisingly, the same result is true for the FEP voltage which is several times larger than the filament voltage. Evidently, turning the loading apparatus off after loading was another critical part of finding a signal.

A Paul-Straubel trap is distinguished from a Paul trap by the absence of endcaps. Usually compensation plates are added on each side of the trap ring electrode along the trap axis. A voltage can be applied to these plates to partially make up for

³That is, just as the laser reaches the first micromotion peak, it jumps several MHz and a sudden decrease in signal is observed. It looks like the transition reached a peak and didn't display micromotion, but, in this case, it has simply skipped over the multiple peaks.

the lack of endcaps. Since the effect of the plate voltage is supposed to reduce the anharmonicities of the trap, the signal should improve. This was not observed but was only attempted once. I found that 30 V was enough to displace the signal out of the laser beam (the signal went to zero), but did not result in ion loss. Once the voltage was removed the signal reappeared. Smaller voltages seemed to have little effect on the signal.

While it was much easier to load a laser cooled ion with the laser initially red detuned by several hundred MHz, I occasionally had success with the laser red detuned only 50 MHz. If I didn't see a signal with the laser detuned by only 50 MHz it was often possible to move it out to 300-400 MHz, scan it back up slowly and have a signal appear. The best explanation for this is that the load is successful either way, but the signal is only visible if the ion loaded happened to be initially cold enough that laser cooling at 50 MHz detuned worked. Hotter (and more typical) ions required the larger initial detuning.

To verify that our power was well above saturation as we expected (see Section 3.2.3), I doubled the laser power to see if there was any change in signal. There was not. In fact, a reduction in power by a factor of 2 left the signal largely unchanged although now a reduction was discernible. In this particular case the UV power out of the enhancement cavity was $50 \mu W$ which meant that about $25 \mu W$ was making it to the trap. A substantial reduction in power from here should still be possible though I didn't try. This is one useful way to increase the signal to background ratio since the background will go down linearly with power.

6.2.9 Alignment Procedure

The following is the procedure used to align and focus the input and output optics. The procedure given is the one that eventually worked. It applies to the twisted wire Paul-Straubel trap which we used to trap single In^+ ions.

- **Input Focus:** The input focus is adjusted by translating one of the input lenses. The correct position of this lens is determined by translating the input lens closest to the trap in a direction perpendicular to the beam. This has the effect of also translating the beam. When the beam is shifted enough that it hits the trap, the image of the spot on the other side of the vacuum chamber

will be eclipsed. Standard focusing techniques can be used with the trap acting as a knife edge.

- **Output Focus:** The (focused) input beam is moved transversely to hit the trap in such a way that it creates a reasonably good UV point source. This UV reflection is imaged on fluorescent paint on the pin hole covering the PMT (the PMT is shifted slightly to one side so that the spot is on the paint instead of the pin hole itself). The PMT is now translated along the output axis until the fluorescent spot is sharply focused. This is now the correct focus *for this spot*—not necessarily the ion. A correction factor was determined but of course it is highly dependent on where the reflected spot comes from. I consistently chose a spot centered in the horizontal direction and shifted to the top edge of the trap. After much experimentation, I found the optimal focus to be 6.8 thousands of an inch closer to the output mirror than the position indicated by the spot focus.
- **Input Vertical Alignment:** The input beam is centered in the trap. There are two orthogonal directions. I arbitrarily called the direction parallel to the output axis the input horizontal direction and orthogonal to that, the input vertical axis. The optimal input vertical position was $16 \mu\text{m}$ towards the pump from center. Center was determined by splitting the distance between eclipsing the beam with two opposite sides of the trap. This correction from center was subject to shifts in the ion position which could easily be caused by contact potentials. Best accuracy was obtained after heating the trap.
- **Input Horizontal Alignment:** The optimal position was $8 \mu\text{m}$ off center moving away from the output mirror.
- **Output Vertical Alignment:** For the output alignment, the PMT is translated in two orthogonal directions perpendicular to the output axis. Here the names vertical and horizontal can be taken literally. The position of the PMT was determined relative to a shadow of the trap on the pin hole created by a lamp placed on the opposite side of the vacuum chamber. If the lamp is placed on the axis joining the trap with the center of the output mirror, then small variations in tilt and position of the lamp will have no effect on the position of

the shadow on the pin hole. The pin hole is then centered in the shadow. The vertical position must account for the correction plate shift (eq. B.3) of about $400 \mu\text{m}$. Consequently, the optimal position was 8.5 thousandths *upward*. This made the pin hole just touch the bottom edge of trap shadow (or the top edge as viewed through a telescope that we used which inverted the image).

- **Output Horizontal Alignment:** Strangely, the output horizontal position required a rather large correction also. I could not account for this, but it was consistently required. Best signal was obtained by moving PMT 3.8 thousandths in the direction away from the pump.

6.2.10 Loading Procedure

The following was the loading/search procedure that finally worked. This procedure is partially automated by the SCAN command on the data collection computer—see Section 4.9.1.

- **Laser Scan Range:** Set the scan range on the laser controller to the desired width (usually 1 GHz in the visible).
- **Initial Laser Detuning:** Set the initial detuning 400 MHz to the red of the transition frequency as determined by the location of Tellurium lines.
- **Filament Voltage:** Set the filament voltage to 48 V (about 4.5 A).
- **Oven Voltage:** Set the oven voltage to 47 V (also about 4.5 A).
- **Pulse Trap RF Off:** Momentarily break the RF connection to the trap to kick out ions (we thought that if too many ions were in the trap from previous loads, they would be hot due to RF heating and the signal would be too broad to be visible).
- **Oven/Filament Warmup:** Turn on the oven and filament such that the filament warms up for 1 minute and the oven warms up for 30 seconds.
- **Pulse E-beam:** Pulse the electron beam on for 2 seconds.

- **Oven Off:** Turn oven voltage off.
- **Filament Off:** Turn filament voltage off.
- **Precool:** After the load is complete leave the laser sitting at its initial detuning for 40 seconds. It is thought that this additional time is required to cool with the weak In^+ cooling transition.
- **Scan Laser:** Scan the laser frequency up towards the ion transition at 1 second per channel. This was often done automatically by having the computer send a ramp voltage signal to the laser controller (the ION command).
- **Set Final Laser Width:** As soon as a signal is detected, detune the laser to the red again and change the scan width to 100 MHz (visible). With the detuning set to 50 MHz the laser is now ready for spectroscopy of this transition.

It can be shown (see Section 2.4.2) that for a hot ion, with a Doppler width of the transition much greater than the secular frequency of the trap ($\Delta_D \gg \omega_i$), that the optimal laser frequency for maximal laser cooling is down-tuned by 1/2 the Doppler width from the transition peak. Of course this position changes as soon as cooling starts to take effect which can be very fast (less than 1 second) for strong transitions. It does, however, suggest a good starting point for setting the laser frequency if hot ions are expected. Since room temperature In^+ ions in a cloud have a Doppler width of 1.5 GHz, we took 800 MHz (in the UV) as a good frequency detuning to start at. This also holds in the “weak binding” regime [Win79a] where now γ (the natural width of the transition) $\gg \omega_i$. In this case, laser cooling has already reduced the Doppler width to be equal to or less than the natural width. So the optimal laser detuning is 1/2 of the *natural* width.

6.3 Conclusions

The ultimate goal of this research is to establish a new frequency standard based on a single laser cooled In^+ ion. The first major step towards this goal has been achieved: the trapping and laser cooling of single In^+ ions. In this thesis I have described the various techniques we used to reach this point: the electrical detection of buffer

gas cooled ion clouds stored in a Paul trap, the optical detection of these clouds, a background suppression scheme that allowed us to improve the SNR, and finally a detailed description of the apparatus and procedures used for loading and detection of laser cooled ions in a Paul-Straubel trap.

6.4 *Where to Next?*

The next goal for this experiment is the observation of quantum jumps *without* a clock laser. This will demonstrate the presence of only one ion in the trap at a time. To reach this goal the SNR needs to be further improved. The first step in this direction is the stabilization of the cooling laser below the natural linewidth of the In^+ cooling transition. This has already been achieved. With an increase in spectral density by a factor of $10\times$ or more it should now be possible to turn the laser power down from $50\ \mu W$ several μW with no significant loss in signal and to achieve a $10\times$ reduction in background (since we are already well into saturation, we don't expect the signal to increase with higher spectral density). An additional benefit is that it should now be possible to tune the laser closer to the peak of the cooling transition without appreciable laser *heating*. This may lead to a larger signal. The apparatus used to stabilize the laser below the In^+ cooling transition width is described in this thesis.

The second step towards quantum jumps is to improve the output optics. An f1.4 Schwarzschild reflecting lens will be used in the future that should double our collection efficiency, reduce reflection losses (our current mirror reflects only 85% whereas the new lens will be coated to transmit better than 99%), and eliminate off axis aberrations, focusing the light more precisely into the PMT pinhole. In addition, the new trap will be oriented such that the optimal output axis will be perpendicular to the output window. This should eliminate astigmatism in the focused output fluorescence which will make the astigmatism correction plate (and its associated losses) unnecessary. The net effect of the new lens will then be at least a factor of 2.5 increase in signal and probably more. Based on calculations in Section 2.7, the stabilized laser taken together with the new lens should provide enough improvements to see quantum jumps.

To build a frequency standard it is necessary to perform high resolution spectroscopy on the narrow clock transition. This will require a second laser system

producing light at 237 nm. Plans are in place for, and development has already started on, a clock laser consisting of a twice doubled 946 nm laser. Since the clock laser has much lower power requirements than the cooling laser, the lower power obtained from laser diodes (compared to dye lasers) is well worth the added stability, reliability and reduction in maintenance. This clock laser system would consist of the diode laser at 946 nm, a KNbO_3 based doubling enhancement cavity producing light at 473 nm, and a BBO based doubling enhancement cavity producing light at 236.5 nm. The latter cavity would be almost identical to the one used currently for the cooling laser. While the ultimate power requirements for the clock laser are probably much less than $1 \mu\text{W}$ at 236.5 nm, higher power may be needed during the initial search for the transition. It is possible that the clock laser system will be able to produce as much as $10 \mu\text{W}$ for this purpose.

Once the clock laser is set up, it must be stabilized. A technique almost identical to the one used for the cooling laser will probably be applied here, with the primary difference being the stability of the external cavity. Greater care must be taken to isolate it thermally and vibrationally since, for frequency standard work, the clock laser needs to have a width of less than 1 Hz. With the clock laser and its supporting apparatus in place, it will be possible to do high resolution spectroscopy. It should also be immediately possible to demonstrate zero point confinement since the In^+ cooling transition width is much less than the secular frequency of the trap.

BIBLIOGRAPHY

- [ARRL] *The ARRL Handbook, 17th Edition*, The American Radio Relay League, 1993.
- [Ash70a] A. Ashkin, *Phys. Rev. Lett.* **24** 156 (1970).
- [Ash70b] A. Ashkin, *Phys. Rev. Lett.* **25** 1321 (1970).
- [Ash78] A. Ashkin, *Phys. Rev. Lett.* **40** 729 (1978).
- [Ber82] J. Bergquist, H. Hemmati and W. M. Itano, *Opt. Commun.* **43** 437 (11/15/82).
- [Ber86] J. Bergquist, R. Hulet, W. Itano, and D. Wineland, *Phys. Rev. Lett.* **57** 1699 (10/6/86).
- [Ber92] J. Bergquist, W. Itano and D. Wineland, *Frontiers in Laser Spectroscopy* (Proc. of the International School of Physics (Enrico Fermi)) edited by T. Hansch and M. Inguscio, North-Holland, Amsterdam, (not published until 1994) p. 359.
- [Ber94] J. Bergquist, private communication (1994).
- [Blo83] L. Bloomfield, *High Resolution Ultraviolet Laser Spectroscopy of Mercury and Helium.*, Ph.D. thesis, Stanford University, 1983 (unpublished).
- [Bor80] M. Born and E. Wolf, *Principles of Optics, Sixth Edition*, Pergamon Press, 1980.
- [Bos89] W. R. Bosenberg, W. S. Pelouch and C. L. Tang, *Appl. Phys. Lett.* **55** 1952 (11/6/89).
- [Boy68] G. D. Boyd and D. A. Kleinman, *J. Appl. Phys.* **39** 3597 (7/68).

- [Bre92] R. Brewer, R. DeVoe and R. Kallenbach, *Phys. Rev. A* **46** R6781 (1992).
- [Bur95] E. Burt and W. Nagourney, to be published.
- [Coh85] C. Cohen-Tannoudji and J. Dalibard, *Europhys. Lett.* **1** 441 (1986).
- [Coo85] R. Cook and H. Kimble, *Phys. Rev. Lett.* **54** 1023 (1985).
- [Cor77] A. Corney, *Atomic and Laser Spectroscopy*, Oxford University Press, 1977.
- [CRC66] *CRC Handbook of Chemistry and Physics, 46th Edition*, The Chemical Rubber Company, 1966.
- [Cut66] L. Cutler and C. Searle, *Proc. of the IEEE* **54** 136 (1966).
- [Deh67] H. Dehmelt, *Advances in Atomic and Molecular Physics*, **3** 53 (1967).
- [Deh73] H. Dehmelt, *Bull. Am. Phys. Soc.* **18** 1521 (1973).
- [Deh75] H. Dehmelt, *Bull. Am. Phys. Soc.* **20** 60 (1975).
- [Deh82] H. Dehmelt, *IEEE Trans. Instrumentations and Measurement* **IM-31** 83 (1982).
- [Dic53] R. Dicke, *Phys. Rev.* **89** 472 (1953).
- [Die89] F. Diedrich, J. Bergquist, W. Itano, and D. Wineland, *Phys. Rev. Lett.* **62** 403 (1/23/89).
- [Dre81] R. W. P. Drever, J. Hall, F. Kowalski, J. Hough, G. Ford, A. Manley and H. Ward, *Appl. Phys. B* **31** 97 (1981).
- [Dre83] R. W. P. Drever, J. Hall, F. Kowalski, J. Hough, G. Ford, A. Manley, and H. Ward, *Appl. Phys. B* **31** 97 (1983).
- [EOMCon] ConOptics, 19 Eagle Road, Danbury, CT 06810.

- [EOMGsa] Gsänger Optoelektronik GmbH; Planegg, Germany.
- [For66] E. N. Fortson, F. Major, and H. Dehmelt, *Phys. Rev. Lett.* **16** 221 (1966).
- [Gar79] F. Gardner, *Phaselock Techniques*, John Wiley & Sons, 1979.
- [Hal81] J. L. Hall, L. Hollberg, T. Baer, and H. Robinson, *Appl. Phys. Lett.* **39** 680 (11/1/81).
- [Han75] T. W. Hansch and A. Schawlow, *Opt. Commun.* **13** 68 (1975).
- [Han80] T. W. Hansch and B. Couillaud, *Opt. Commun.* **35** 441 (12/80).
- [Hil87] D. Hills and J. L. Hall, *Rev. Sci. Instrum.* **58** 1406 (8/87).
- [Hor93] P. Horowitz and W. Hill, *The Art of Electronics, 2nd Edition*, Cambridge University Press, 1993.
- [Jan84] G. Janik, Ph. D. Thesis, University of Washington, 1984 (unpublished).
- [Jan85] G. Janik, W. Nagourney, and H. Dehmelt, *J. Opt. Soc. Am. B* **2** 1251 (8/85).
- [Jav86] J. Javanainen, *Phys. Rev. A* **33** 2121 (1986).
- [Kat86] K. Kato, *IEEE J. Quantum Electron.* **22** 1013 (1986).
- [Kop58] H. Kopfermann, *Nuclear Moments*, Academic Press (1958).
- [Lar93] P. Larkins and P. Hannaford, *Z. Phys. D* **27** 313 (1993).
- [Lan76] L. Landau and E. Lifshitz, *Mechanics, Third Edition, Course of Theoretical Physics Volume 1*, Pergamon Press 1976.
- [Let77] V. Letokhov, *Comments At. Mol. Phys.* **6** 119 (1977).
- [Lou83] R. Loudon, *The Quantum Theory of Light, Second Edition*. Oxford University Press, 1983.

- [Mci87] D. H. McIntyre, *High Resolution Laser Spectroscopy of Tellurium and Hydrogen: A Measurement of the Rydberg Constant*, Ph. D. thesis, Stanford University, 1987 (unpublished).
- [Nag83] W. Nagourney, G. Janik, and H. Dehmelt, *Proc. Nat. Acad. Sci.* **80** 643 (1983).
- [Nag86] W. Nagourney, J. Sandberg, and H. Dehmelt, *Phys. Rev. Lett.* **56** 2797, (1986).
- [Nag90] W. Nagourney, N. Yu, and H. Dehmelt, *Opt. Commun.* **79** 176 (10/15/90).
- [Neu78] W. Neuhauser, M. Hohenstatt, P. Toschek and H. Dehmelt, *Phys. Rev. Lett.* **41** 233 (7/24/78).
- [Neu80] W. Neuhauser, M. Hohenstatt, P. Toschek and H. Dehmelt, *Phys. Rev. A* **22** 1137 (1980).
- [Pas38] F. Paschen and J. Campbell, *Ann. Phys.* **31** 29 (1938).
- [Pau58] W. Paul, O. Osberghaus, and E. Fischer, *Forshungsber. Wirtsch. Verkehrsministeriums Nordrhein-Westfalen* No. 415.
- [Pei94] E. Peik, G. Holleman and H. Walther, *Phys. Rev. A* **49** 402 (1994).
- [Pre89] J. Prestage, G. Dick, and L. Malecki, *J. Appl. Phys.* **66** 1013 (1989).
- [Pre92] W. Press, *Numerical Recipes in C, Second Edition*, Cambridge Press (1992).
- [Rag89] P. Raghavan, *At. Data Nucl. Data Tables* **42** 189 (1989).
- [Rai92] M. Raizen, J. Gilligan, J. Bergquist, W. Itano and D. Wineland, *J. Mod. Opt.* **39** 233 (1992).
- [Ram53] N. Ramsey, *Nuclear Moments*, John Wiley & Sons (1953).
- [Rau94] M. Rauner, Master's Thesis, University of Colorado, 1994 (unpublished).

- [Rei65] F. Reif, *Fundamentals of Statistical and Thermal Physics*, McGraw-Hill (1965).
- [Ric57] M. Rice and R. Pound, *Phys. Rev.* **106** 953 (1957).
- [San93] J. Sandberg, Ph. D. Thesis, Massachusetts Institute of Technology, 1993 (unpublished).
- [Sau86] Th. Sauter, W. Neuhauser, R. Blatt and P. Toschek, *Phys. Rev. Lett.* **57** 1696 (10/6/86).
- [Sch86] A. Schenzle, R. DeVoe, and R. Brewer, *Phys. Rev. A* **33** 2127 (1986).
- [Sch93] C. Schrama, E. Peik, W. Smith and H. Walther, *Opt. Commun.* **101** 32 (1993).
- [Ste95] A. Steinbach, M. Rauner, and J. Bergquist, to be published.
- [Sto76] D. Stout and M. Kaufman, *Handbook of Op Amp Circuit Design*, McGraw-Hill, 1976.
- [Str55] H. Straubel, *Naturwissenschaften* **18** 506 (1955).
- [Sve89] O. Svelto, *Principles of Lasers*, Plenum, 1989.
- [Van86] J. Van de Vegte, *Feedback Control Systems*, Prentice Hall, 1986.
- [Wak92] I. Waki, S. Kassner, G. Birkl, and H. Walther, *Phys Rev. Lett.* **68** 2007 (1992).
- [Win75] D. Wineland and H. Dehmelt, *Bull. Am. Phys. Soc.* **20** 637 (1975).
- [Win78] D. Wineland, R. Drullinger, and R. Walls, *Phys. Rev. Lett.* **40** 1639 (1978).
- [Win79a] D. Wineland and W. Itano, *Phys. Rev. A* **20** 1521 (10/79).
- [Win79b] D. Wineland, *J. Appl. Phys.* **50** 2528 (1979).

- [Win81] D. Wineland, and W. Itano, *Phys. Lett. A* **82A** 75 (3/9/81).
- [Win87] D. Wineland, W. Itano, J. Bergquist and R. Hulet, *Phys. Rev. A* **36** 2220 (9/1/87).
- [Woo80] G. Woodgate, *Elementary Atomic Structure*, Oxford University Press (1980).
- [Yar89] A. Yariv, *Quantum Electronics, Third Edition*, John Wiley and Sons (1989).
- [Yu90] N. Yu, W. Nagourney, and H. Dehmelt, *J. Appl. Phys.* **69** 3779 (3/15/90).
- [Yu91] N. Yu, W. Nagourney, and H. Dehmelt, *J. Appl. Phys.* **69** 3779 (3/15/91).
- [Yu95] N. Yu and W. Nagourney, *J. Appl. Phys.* **77** 3623 (4/15/95).
- [Yu95a] N. Yu, private communication, 1995.
- [Zha95] X. Zhao, N. Yu, H. Dehmelt, and W. Nagourney, *Phys. Rev. A* **51** 4483 (6/95).

Appendix A

INDIUM ION SUMMARY

This Appendix summarizes some of the aspects of Indium II (singly ionized Indium, usually referred to as In^+) that are relevant to this experiment. Table A.1 gives information on the cooling and clock transitions. There are two stable isotopes of Indium: $A = 113$ and $A = 115$. Information found in [Rag89] (except where noted) on abundances, magnetic moments, and quadrupole moments is compiled in Table A.2.

Table A.3 summarizes the saturation laser power for each transition and for various widths. Note that 15 GHz corresponds to the Doppler width associated with roughly 1/10 the well depth of our Paul trap, 1.5 GHz corresponds to the Doppler width at 300 K, and $\Delta\omega_D = 0$ indicates that the Doppler width is much smaller than the other widths. For the clock transition, we are only interested in the $\Delta\omega_D = 0$ case because observations on this transition will only take place when the ion is cold. See Section 3.2.3 for more details.

Table A.4 summarizes the estimated fluorescence signal from the cooling transition that should be observed as a function of number of ions, laser power, laser width and Doppler width. An ion cloud cannot be buffer gas cooled below 300 K so the last two entries in the first row are not applicable. Similarly, the fluorescence signal from a single ion is not interesting *unless* it is cooled well below 300 K. Note that these values do not take into account the effect of background suppression. Finally, the

Table A.1: Frequencies and Lifetimes for Indium II Transitions.

transition:	$5s^2 \ ^1S_0 - 5s5p \ ^3P_1$	$5s^2 \ ^1S_0 - 5s5p \ ^3P_0$
wavelength:	230.6 nm	236.5 nm
$\nu(cm^{-1})$ [Lar93]:	43,351.614(2)(5) (F=11/2) 43,350.351(3)(6) (F=9/2) 43,349.289(3)(6) (F=7/2)	42,275.986(2)(7) (F=9/2)
τ [Pei94]:	0.44 μs	0.14 s

Table A.2: Data on abundances, magnetic moments, and quadrupole moments for each stable isotope. This information was found in [Rag89] except where noted.

$$\begin{aligned}
 \text{abundance}_{113} &= 4.3\% \\
 \text{abundance}_{115} &= 95.7\% \\
 \tau_{113} &= \text{stable} \\
 \tau_{115} &= 4.41 \times 10^{14} \text{y.} \\
 I_{113} &= I_{115} = +9/2 \\
 \mu_I^{113} &= 5.5289(2) \mu_N \text{ (ref. [Ric57])} \\
 \mu_I^{115} &= 5.5408(2) \mu_N \\
 Q_{113} &= +0.799 \\
 Q_{115} &= +0.810
 \end{aligned}$$

Table A.3: Saturation laser power for the In^+ transitions with various combinations of widths: ω_L corresponds to the laser width and ω_D corresponds to the Doppler width of the transition.

transition:	cooling			clock
	$\Delta\omega_D = 15GHz$	$\Delta\omega_D = 1.5GHz$	$\Delta\omega_D = 0$	$\Delta\omega_D = 0$
$\Delta\omega_L = 5MHz$	4 mW	402 μW	1.34 μW	1.24 μW
$\Delta\omega_L = 30kHz$	4 mW	402 μW	97 nW	7.4 nW

Table A.4: Estimated fluorescence signal for the In^+ cooling transition.

	$\Delta\omega_D = 1.5GHz$ $\Delta\omega_L = 5MHz$	$\Delta\omega_D = 0$ $\Delta\omega_L = 5MHz$	$\Delta\omega_D = 0$ $\Delta\omega_L = 30kHz$
$N = 50,000$ $P = 50 \mu W$	$13,000 s^{-1}$	NA	NA
$N = 1$ $P = 1 \mu W$	NA	$661 s^{-1}$	$914 s^{-1}$

saturation power for a laser cooled ion irradiated by a 30 kHz laser is about 100 nW. The power for the last entry in the second row has been reduced to this value. See Section 3.2.2 for more details.

Appendix B

SOME USEFUL FORMULAE

This appendix summarizes several useful formulae found in the text. A derivation of the displacement of the output fluorescence due to the presence of the astigmatism correction plate (see Section 4.10.2) is derived in Section B.1. A discussion of the Gaussian lineshape used to fit Doppler broadened spectra is given in Section B.2. Finally a list of other formulae used in the text is given in Section B.3 along with references to where they are either derived, or at least discussed.

B.1 Correction Plate Displacement

This section gives the derivation of the displacement between two beams of different wavelengths incident on the correction plate used to reduce astigmatism in the ion fluorescence. Further discussion of this plate may be found in Section 4.10.2. If two collinear beams of different wavelengths are incident at an angle θ_i on a material with indices of refraction n_1 and n_2 (for the two wavelengths) and thickness t , then they will emerge on the other side of the material still collinear but displaced from each other. See Figure B.1. The two angles of refraction through the plate are θ_1 and θ_2 . The respective displacements of each beam from the original direction, measured along the output face of the plate are d_1 and d_2 . The difference between these two displacements is Δd and the component of this difference transverse to the output propagation direction is Δd_t . From simple geometry we have

$$\Delta d_t = \Delta d \cos \theta_i \quad (\text{B.1})$$

where Δd can be written as

$$\Delta d = t(\tan \theta_i - \tan \theta_1 - \tan \theta_i + \tan \theta_2). \quad (\text{B.2})$$

Combining equations B.1 and B.2 and applying Snell's law, one obtains:

$$\Delta d_t = t(\tan(\sin^{-1}(\frac{\sin \theta_i}{n_1})) - \tan(\sin^{-1}(\frac{\sin \theta_i}{n_2}))) \cos \theta_i. \quad (\text{B.3})$$




Figure B.1: Astigmatism correction plate geometry.

In our situation there were two cases to consider. First, the output fluorescence goes through the output window of the vacuum chamber. This window has thickness $t_w = 5\text{mm}$ and incident angle $\theta_i^w = 20^\circ$. The second case is the correction plate with thickness $t_c = 9.53\text{mm}$ and incident angle $\theta_i^c = 53^\circ$. The first wavelength is 231 nm for our cooling transition. The second wavelength is taken to be 550 nm, or in the middle of the visible spectrum. Both materials are quartz so that $n_1 = 1.52$ and $n_2 = 1.46$. Substituting these values into eq. B.3 gives a transverse displacement due to the window, $\Delta d_w \approx 47\mu\text{m}$, and due to the correction plate, $\Delta d_c \approx 343\mu\text{m}$ respectively. The fluorescence-detecting PMT is positioned just after the plate. So the displacement along the PMT aperture and perpendicular to the axis of the fluorescent beam is

$$\Delta d_{c+w} \approx 390\mu\text{m}. \quad (\text{B.4})$$

B.2 Gaussian Fits

This section describes the Gaussian lineshape used to fit Doppler broadened data. The true lineshape is usually a convolution of Gaussian and Lorentzian lineshapes, called a Voigt profile. At the temperatures found in a buffer gas cooled ion cloud, however, the Lorentzian contribution is quite small by comparison to the Gaussian contribution. Therefore, the true lineshape can be closely approximated by a Gaussian. The equation used for Gaussian fitting was:

$$t = k_0 + k_1 e^{-((x-k_2)/k_3)^2} \quad (\text{B.5})$$

where t is the number of transitions per channel, k_i are the fit parameters, and x is the channel number. The values k_0 and k_1 are in units of counts per integration channel. The values k_2 and k_3 are in number of channels. From k_3 , which is the $1/c$ half width, we obtain the full width at half maximum (FWHM) in MHz from the equation:

$$\delta_{f_{whm}} = 2\sqrt{\ln 2} k_3 \Delta f = 1.67 k_3 \Delta f \quad (\text{B.6})$$

where Δf is the width of each channel in MHz. Finally, one can obtain an approximate temperature from this:

$$T = (\delta_{f_{whm}}/1500)^2 300 \quad (\text{B.7})$$

where room temperature, 300 K, corresponds to a FWHM of 1500 MHz for In^+ [Cor77, Section 8.3].

B.3 Summary of Useful Equations

What follows is a listing of various useful equations found throughout the text. A reference to the appropriate location in the text for each equation is also given.

Buffer Gas Collision Rate (eq. 2.34):

$$R = \frac{4\pi r^2 P_H}{kT_H} \sqrt{\frac{8k}{\pi} \left(\frac{T_I}{M} + \frac{T_H}{m} \right)}$$

Energy Transfer From In^+ To Buffer Gas Per Collision (eq. 2.29):

$$\frac{\delta E}{E} \simeq 0.140$$

Doppler Cooling Limit (eq. 2.40):

$$T_{min}^{wb} = \frac{\hbar\gamma}{2k}$$

Sideband Cooling Limit (eq. 2.45):

$$\langle n_z \rangle_{min}^{SB} = 5/16(\gamma/\omega_z)^2$$

Sideband Cooling Rate (eq. 2.43):

$$\frac{dE}{dt} = \hbar m \omega_z \frac{I \sigma_0}{\hbar \omega_0} J_m^2(kx_-)$$

Power Broadening (eq. 2.62):

$$\Delta\nu_{PB} = 2\gamma \sqrt{1 + \frac{3\lambda^3 \tau P}{2\pi^2 c \hbar A}}$$

Secular Motion Amplitude (eq. 2.46):

$$x_a = \sqrt{\frac{2\langle n_z \rangle \hbar}{M(2\pi)^2 \nu_z}}$$

Paul Trap Capacity (eq. 2.38):

$$N_{max} = 2.57 \times 10^5 \bar{D} z_0$$

Paul Trap Secular Frequency (eq. 2.20):

$$\omega_z = \frac{4\sqrt{2}eV_0}{m\Omega r_0^2} \quad (r_0 = 2z_0)$$

Paul Trap Well Depth (eq. 2.14):

$$\bar{D} = \frac{eA_0^2 r_0^2}{m\Omega^2} \quad (r_0 = 2z_0)$$

Maximum Fluorescence Count Rate (eq. 2.56):

$$R_{max} = I_0 \frac{2\sqrt{\ln 2} \lambda^3}{2\pi^{1/2} \hbar c} \frac{\gamma}{\sqrt{\gamma^2 + \Delta\omega_D^2 + \Delta\omega_L^2}}$$

Saturation Laser Power (eq. 2.59):

$$P_{sat} = \frac{\pi^{3/2} \hbar c \omega_0^2 \sqrt{\gamma^2 + \Delta\omega_D^2 + \Delta\omega_L^2}}{\sqrt{\ln 2} \lambda^3}$$

Gaussian Beam Spot Size (eq. 4.8):

$$\omega_0 = \frac{\Delta x_{10-90}}{1.28}$$

Gaussian Lineshape (eq. B.5):

$$t = k_0 + k_1 e^{-((x-k_2)/k_3)^2}$$

Correction Plate Displacement (eq. B.3):

$$d = t \left(\tan\left(\sin^{-1}\left(\frac{\sin\theta_i}{n_1}\right)\right) - \tan\left(\sin^{-1}\left(\frac{\sin\theta_i}{n_2}\right)\right) \right) \cos\theta_i$$

Appendix C

ELECTRICAL CIRCUITS

This appendix contains schematics for several of the key pieces of electronic equipment used in the experiment.

C.1 Trap RF Drive Amplifier

Figure C.1 shows the circuit used to amplify the RF drive (Ω) on the trap ring electrode. The input was usually several volts at about 10 MHz. The VFET gives about 20 dB of gain. The output is routed to a step-up coil which increases the voltage amplitude of the signal by about another factor of 10.

C.2 Electrical Detection Preamp

Figure C.2 shows the schematic for the preamp used in the electrical detection of ion clouds. The circuit is based on one described in [Hor93, p. 428]. It uses two followers on the input to give a large input impedance and better common mode rejection ratio (CMRR). Good CMRR (100 dB or more) is a requirement since the signal is still together with other (large) sources of RF at this point.

C.3 Short Term Laser Stabilization Electronics

This section describes the schematic for the short term laser stabilization electronics. This circuit is an adaptation of one found J. Sandberg's Ph.D. thesis [San93].

Figure C.3 shows the electronics used to detect the error signal for laser stabilization. The FND-100Q is the photodiode detector. The error signal is fed into the CLC425 wide band, low noise amplifier and then into the Mini-Circuits MAR-3 power amplifier. The signal is both AC and DC coupled to the output. The AC coupled output is generally used as the input to the RF mixer in the next stage.

Figure C.4 gives the schematic for the reference oscillator and the RF mixer. The 24 MHz crystal oscillator determines the frequency of the modulation voltage applied




Figure C.1: Amplifier circuit used to boost the trap RF drive.




Figure C.2: Preamp circuit used for electrical detection of an ion cloud.

to the EOM to create sidebands on the laser. The output of the oscillator at J2 is routed to the EOM via a Motorola CA2832C power amplifier (not shown) and also to the LO input of the Mini-Circuits SRA-1W mixer via a phase shifter (this is not shown either—it could be as simple as just a length of cable). The RF error signal from the detector is routed to the RF port of the mixer. The DC error signal is then available on the IF port of the mixer.

Figure C.5 shows the integrator network. The DC error signal from the mixer is connected to J1 on the non-inverting input of the wide band integrator (CLC425). The output of this integrator is connected to one input of the intra-cavity EOM, to the slow path transfer electronics, and to the inverting input of the HV integrator. The signal at B enables or disables the HV integrator. The output of the HV integrator, at J3, is connected to the second input of the intra-cavity EOM. When the laser is locked to the external cavity, the signal from both integrators is being applied to the EOM. When the laser is not locked, the HV integrator is disabled by LOCK CTL. In this way, the ringing and overshoot that may be associated with a high gain 2-pole rolloff is avoided. The high gain of both cascaded integrators is only applied when lock has been acquired.

Figure C.6 shows the cross over network used to match the gain of the intra-cavity EOM to that of the tweeter at the frequency where one takes over from the other. This is called the slow path transfer circuit because the inherent 699 transducers have much lower bandwidth than the intra-cavity EOM. The match is accomplished by tuning the Slow Path Gain Adj. pot. The signal at B from the Lock CTL circuitry enables or disables the integrators. The output at C goes to the redundant locking circuitry which uses the 699-21 servo loop as a backup for the high bandwidth EOM-based servo loop.

Figure C.7 shows the redundant locking circuit. The LOCK DET IN input comes from a photodiode that detects leakage light coming out of the reference cavity. When the laser is locked the leakage light will be more intense than when the laser is out of lock. The lock signal is routed through an LED which acts as a lock indicator. The output at B is the lock control signal used by other parts of the circuit and represents the lock detect signal unless that is overridden by the Master Lock Control switch.

The error signal from the slow path transfer circuit is routed to C. The error signal from the 699-21 (derived from its own cavity) is routed to 699 INT ERR IN at J6. If

the laser is locked to the external cavity then the error signal at C will be enabled and routed to the 699 INT ERR OUT port at J5. There it is routed to the existing 699 transducers as if this were its very own error signal. The 699-21 circuitry is broken between TP1 and R17 on board 1A9 and routed to its backpanel. The output is then routed to J6 on our circuit with J5 on our circuit then being routed back to the input on the 699-21 backpanel. If LOCK CTL indicates OFF (not locked to the external cavity) then the error signal at C is disabled and instead the 699 error signal is routed directly back into the 699-21. If LOCK CTL indicates ON then the 699 error signal is prevented from going back to the 699-21; instead it is routed to the 699 EXT OUT at J7. From here it is connected to the 699 external control.

In summary, the redundant locking circuit drives the 699 transducers with either the error signal from the external cavity, or from the 699 cavity, depending on whether the laser is locked to the external cavity. It also keeps the 699 cavity close to the external cavity by driving the 699 external frequency control. This makes it possible for the circuit to quickly switch between error signals if the high frequency servo loop loses lock.

C.4 Drift Stabilization Electronics

The schematic for the drift stabilization electronics was not available on-line at the time of this writing, so I have included a block diagram of it in Figure C.8 instead. The 4 MHz oscillator drives the EOM that creates the sidebands on the laser. Mixing this with the detected saturation signal produces a DC error signal on the IF port connected to the input of the phase sensitive detector (PSD). The reference for the PSD comes from the 10 kHz chopping signal which is applied, together with 94 MHz shift signal, to the AOM. The 10 kHz portion of this signal chops the saturation beam on and off and the 94 MHz portion causes the AOM to create sidebands at ± 47 MHz. This modulation frequency, together with which sideband is chosen, determines the frequency shift between the fundamental laser frequency and the Te saturation line. The output of the PSD is then routed into an integrator and from there to either the 699-21 external frequency control or the external cavity used for short term stabilization. The former is used if the latter is not operational.




Figure C.3: Laser stabilization detection electronics.




Figure C.4: Reference oscillator electronics for laser stabilization.




Figure C.5: Integration electronics for laser stabilization.




Figure C.6: Slow path transfer electronics for laser stabilization.




Figure C.7: Redundant locking electronics for laser stabilization.




Figure C.8: Block diagram for the drift stabilization electronics.

C.5 Enhancement Cavity Locking Electronics

The locking electronics for the SHG enhancement cavity is used to lock the cavity to the fundamental laser frequency. This makes it possible to tune the laser and, without losing lock, to tune the second harmonic. The circuit is described in more detail in Section 4.6.4. Figure C.9 shows the detection and integration stages. The detection stage consists of the two photodiodes that detect the orthogonally polarized beams, followed by an amplifier and a unity gain inverter. Both the error signal and the inverted error signal are available. The integration stage consists of a follower for monitoring purposes, and three integrators. Each integrator can be reset by the HV rail detect circuitry which determines whether the error signal has gone out of the HV amplifier range. The second and third integrators are prevented from contributing immediately to the gain by an RC circuit. Maximum gain is applied only when the circuit has acquired lock.

Figure C.10 shows the HV and HV rail detect stages. The HV stage takes its input either from the integrators or from the sweep signal. The latter allows sweeping the cavity across the laser line; particularly useful during alignment. The output of the HV stage goes to the cavity PZT. The HV rail detect circuitry determines whether the HV amplifier has reached one of its rails and so is not able to generate a valid error signal any longer. If this is the case, a reset signal is sent to the integrators. The cavity is momentarily taken out of lock while the electronics attempts to reacquire lock on a closer cavity fringe.

Figure C.11 shows the sweep circuit that generates the internal ramp voltage used for sweeping the cavity across the laser line.

C.6 Gating and Detection Electronics

Figures C.12 through C.14 show the circuitry for the background suppression gating electronics. This circuit was designed by Bryan Venema and Phil Williams of our electronics shop. With it we are able to control the duty cycle and frequency of the laser light modulation and the duty cycle and relative phase of the detection gate. The detected photon counts are input at J2 and gated by the signal at C to give a gated count signal at J3.




Figure C.9: Schematic of the locking electronics for the SHG enhancement cavity—part 1. This schematic shows the detection and integration stages.




Figure C.10: Schematic of the locking electronics for the SHG enhancement cavity—part 2. This schematic shows the HV stage and the HV rail detector.




Figure C.11: Schematic of the locking electronics for the SHG enhancement cavity—part 3. This schematic shows the circuit used to generate a ramp voltage for the internal sweep mechanism.




Figure C.12: Background suppression gating electronics—part 1.




Figure C.13: Background suppression gating electronics—part 2.




Figure C.14: Background suppression gating electronics—part 3.

VITA

Eric Alban Burt, son of George James and Dierdre Kisich Burt, was born on March 27, 1957 in Berkeley, California. He majored in math at Ann Arbor Pioneer High School in Ann Arbor, Michigan, graduating in the top 7% of his class in 1975. He went on to major in math and minor in physics at the University of Michigan where he graduated with honors and a Bachelor of Science degree in 1979. After receiving his undergraduate degree, Eric worked for 6 years as a systems programmer at Ford Motor Company, and then for 2 years as a research scientist specializing in the development of image processing software at the Environmental Research Institute of Michigan (ERIM). During his time at ERIM, he returned to school part time to study physics. In 1988 he studied physics full time at the University of Washington in Seattle, and in 1989 was admitted to the graduate program. He received his Master of Science degree in 1990 and in 1991 joined H. Dehmelt's and W. Nagourney's single ion research lab. In 1995, Eric received the Doctor of Philosophy degree for his work on trapped laser cooled single Indium ions. Eric also maintains a strong interest in sports, having attended the University of Michigan on a track and field athletic scholarship and winning a medal at the 1991 United States National Cycling Championships.